

The Orbit Problem in Higher Dimensions

Ventsislav Chonev, Joël Ouaknine, James Worrell

Oxford University Department of Computer Science

Abstract. Our work focuses on the extended orbit problem: determining whether a target vector space V may be reached from a point x under repeated application of a linear transformation A . We give a polynomial-time algorithm for the case when V is one-dimensional. We also show that when V is two- or three-dimensional, the problem is in the complexity class NP^{EqSLP} , that is, NP with an oracle for deciding zeroness of arithmetic circuits, and hence in NP^{RP} .

1 Introduction

The *orbit problem* is defined as follows:

Given a rational matrix A and rational vectors x, y , decide whether there exists a non-negative integer n such that $A^n x = y$.

It may be interpreted as a reachability problem from point to point in a linear system with transitions described by A . In a well-known paper [1], Kannan and Lipton showed a polynomial-time algorithm for the problem. The algorithm first reduces the orbit problem to the matrix power problem:

Given a rational matrix A and a polynomial $p \in \mathbb{Q}[x]$, decide whether there exists a non-negative integer n such that $A^n = p(A)$.

Then this problem is solved using the correspondence

$$A^n = p(A) \Rightarrow \alpha^n = p(\alpha) \text{ for all eigenvalues } \alpha \text{ of } A$$

Kannan and Lipton proved that if A has some eigenvalue which is not a root of unity, a polynomial bound m exists such that if $A^n = p(A)$, then $n < m$. With some effort they extended the decision method to cover matrices whose eigenvalues are all roots of unity, also in polynomial time.

In their conclusion, they considered the following *extended orbit problem*:

Given a rational matrix A and a vector subspace V specified by a basis of rational vectors, decide whether there exists a non-negative integer n such that $A^n x \in V$.

The extended orbit problem replaces the point y with a target space V . Kannan and Lipton conjectured that a polynomial-time algorithm exists for the case where $\dim(V) = 1$, and suggested that decidability be investigated in the cases

$\dim(V) = 2$ and $\dim(V) = 3$. Later work by Arvind and Vijayaraghavan [2] placed the original orbit problem in the logspace counting hierarchy GapLH. To the best of our knowledge, there has been no progress on the decidability of the extended orbit problem or its fixed-dimension cases.

A related problem is the *inhomogeneous extended orbit problem*, where the target is an affine space $V + z$. The case $\dim(V) = 0$ is Kannan and Lipton's original problem. In general, the k -dimensional inhomogeneous version reduces to the $(k + 1)$ -dimensional homogeneous version with the addition of a dummy component to all vectors.

Another related problem is the *chamber hitting problem*. It replaces the target space with an intersection of half-spaces. In reference [3], the chamber hitting problem is related to the $P_{\mathbb{B}}$ -realizability problem on regular languages: determine whether the language of a given finite deterministic automaton contains a word of the form $\#w_1\#w_2\#\dots\#w_N\#$, where $N = 2^n$ is a fixed power of 2 and the words w_1, \dots, w_N are all the binary words of length n in some order.

The orbit problem is also closely linked to *Skolem's problem*. Given a linear recurrent sequence described by its initial conditions and a recurrence formula, the problem is to decide whether 0 is an element of the sequence. An equivalent formulation is to determine whether there exists a non-negative integer n such that $x^T A^n y = 0$ for given vectors x, y and a square matrix A . The equivalence of these formulations is easy to show using the companion matrix of the sequence and the Cayley-Hamilton theorem. The dimension of A in the second formulation equals to the depth of the recurrence in the first.

The decidability of Skolem's problem is open. A well-known result, the Skolem-Mahler-Lech theorem, states that the set of zeroes of any linear recurrence is the union of a finite set and finitely many arithmetic progressions [4,5,6,7]. Moreover, it is known how to effectively compute the arithmetic progressions [7]. The main difficulty to deciding Skolem's problem is to determine whether the finite component of the set of witnesses is empty.

Skolem's problem is trivial for recurrences of depth 1. A proof of decidability for depth 2 may be found in [8]. The problem was open for depth 3 and 4 until Vereshchagin [9] and Mignotte, Shorey and Tijdeman [10] proved it decidable in 1985 using Baker's results for linear forms in logarithms [11]. The case of depth 5 was considered in [12]. Skolem's problem for arbitrary depth was shown NP-hard by Blondel and Portier [13].

Kannan and Lipton [1] pointed out that Skolem's problem is easy to reduce to the orbit problem. Deciding whether $x^T A^n y = 0$ for some n is equivalent to deciding whether $A^n y$ is in $(\text{span}\{x\})^\perp$ for some n . A reduction also exists in the other direction. Let (A, x, V) be an instance of the orbit problem, where V has basis $\{v_1, \dots, v_k\}$. Let $\{v_{k+1}, \dots, v_m\}$ be a basis for V^\perp . Then we have

$$A^n x \in V \iff \bigwedge_{i=k+1}^m v_i^T A^n x = 0$$

Given a yes-no oracle for Skolem's problem, the set of witnesses of a Skolem instance may be effectively computed. This is done by extracting the periodic sets

of zeroes and then iterating the sequence, using the oracle at each step to determine if all the zeroes making up the finite component have been found. So we calculate the sets of witnesses for each instance and check whether the intersection is non-empty. Or alternatively, if s_i are the linear recurrences corresponding to the Skolem instances $v_i^T A^n x = 0$, we can define the sequence

$$S(j) = \sum_{i=k+1}^m (s_i(j))^2$$

It is classical that linear recurrences are closed under addition and termwise multiplication, so S is a linear recurrence whose zeroes are exactly the witnesses to the orbit problem instance.

Observe that in this reduction, the resulting Skolem instances have depth dependent on the size of the matrix in the orbit instance, not on the dimension of the target space. Since we are unable to solve Skolem instances for depth greater than 4, this severely limits the usefulness of the reduction.

In this paper, we show how to solve the orbit problem with a matrix of arbitrary size and target space of dimension k via Skolem's problem for depth $k + 1$. In particular, we show that the one-dimensional orbit problem is decidable in polynomial time and that the two- and three-dimensional versions are in the complexity class NP^{EqSLP} and hence in NP^{RP} . We first use a reduction to higher-dimensional versions of the matrix power problem, in the style of Kannan and Lipton. Then we construct an equivalent system of equations from the eigenvalues of the given matrix. From the system, we extract tuples of equations and then apply known techniques from Skolem's problem to bound the exponent n .

The structure of the paper is as follows. In section 2 we reduce the fixed-dimension orbit problem to the fixed-dimension matrix power problem. In section 3, we prove that the one-dimensional version is decidable in polynomial time. In sections 4 and 5 we show decidability for the two- and three-dimensional cases. Finally, section 6 concludes the paper. Appendix A briefly outlines how algebraic numbers may be represented and used in arithmetic. Appendix B states some basic results from algebraic number theory. Appendix C states Baker's theorem, van der Poorten's theorem and some applications relevant to both the orbit problem and Skolem's problem. Appendices D, E and F give all the technical lemmas necessary for proving decidability of Skolem's problem in depth up to 4.

2 Reduction

2.1 Matrix power problem

Suppose we are given a rational matrix A , a rational vector x and a target vector space V specified by a basis of rational vectors y_1, \dots, y_k . We wish to decide whether there exists $n \in \mathbb{N}$ such that $A^n x \in V$.

Observe we can rescale A in polynomial time by the least common multiple of all denominators appearing in A . This reduces the problem to a version of itself

but with the additional assumption that A is an integer matrix. This ensures the eigenvalues of A are all algebraic *integers*.

Calculate

$$\nu = \max \{m \mid x, Ax, \dots, A^m x \text{ are linearly independent}\}$$

Let $\mathcal{B} = \{x, Ax, \dots, A^\nu x\}$, $U = \text{span}(\mathcal{B})$ and $D = [x \ Ax \ \dots \ A^\nu x]$. It is clear that U is invariant under the linear transformation A , so consider the restriction of A to U . If $[b_0, \dots, b_\nu]^T$ are the coordinates of $A^{\nu+1}x$ with respect to \mathcal{B} , then this restriction is described by the matrix

$$M = \begin{bmatrix} 0 & 0 & \dots & 0 & b_0 \\ 1 & 0 & \dots & 0 & b_1 \\ 0 & 1 & \dots & 0 & b_2 \\ & & \ddots & & \vdots \\ 0 & 0 & \dots & 1 & b_\nu \end{bmatrix}$$

That is, if $z = Dz_{\mathcal{B}}$, then $Az = DMz_{\mathcal{B}}$. By induction, for all $n \in \mathbb{N}$, $A^n x = DM^n x_{\mathcal{B}}$, where $x_{\mathcal{B}} = [1, 0, \dots, 0]^T$. Next we calculate a basis for $W = U \cap V$, let this basis be $\{w_1, \dots, w_t\}$ and let $w_i = Dw'_i$ for all i . Now,

$$A^n x \in V \iff A^n x \in W \iff M^n x_{\mathcal{B}} \in \text{span}\{w'_1, \dots, w'_t\}$$

Thus, the problem is reduced to a version of itself where the underlying matrix is *invertible*. Notice also that the matrix M describes a restriction of the linear transformation A , so its eigenvalues are a subset of the eigenvalues of A . In particular, since A was rescaled to an integer matrix, the eigenvalues of M are algebraic integers as well.

Define the matrices T_1, \dots, T_t by

$$T_i = [w'_i \ Mw'_i \ \dots \ M^\nu w'_i]$$

We will show that $M^n x_{\mathcal{B}} \in \text{span}\{w'_1, \dots, w'_t\}$ if and only if $M^n \in \text{span}\{T_1, \dots, T_t\}$. If for some coefficients a_i we have

$$M^n = \sum_{i=0}^t a_i T_i$$

then considering the first column of both sides, we have

$$M^n x_{\mathcal{B}} = \sum_{i=0}^t a_i w'_i$$

Conversely, suppose $M^n x_{\mathcal{B}} = \sum_{i=0}^t a_i w'_i$. Then note that $x_{\mathcal{B}}, Mx_{\mathcal{B}}, \dots, M^\nu x_{\mathcal{B}}$ are just the unit vectors of size $\nu + 1$. Multiplying by M^j for $j = 0, \dots, \nu$ gives $M^{n+j} x_{\mathcal{B}} = \sum_{i=0}^t a_i M^j w'_i$. The left-hand side is exactly the $(j + 1)$ -th column of

M^n , whereas $M^j w'_i$ on the right-hand side is exactly the $(j+1)$ -th column of T_i . So we have $M^n = \sum_{i=0}^t a_i T_i$.

Thus, we have reduced the orbit problem to the *matrix power problem*: determining whether some power of a given matrix lies inside a given vector space of matrices. Now we will perform a further reduction step. It is clear that out of the space $T = \text{span}\{T_1, \dots, T_t\}$ it suffices to consider only matrices of the shape $p(M)$ where p is a polynomial. We find a basis for the space $P = \{p(M) \mid p \in \mathbb{Q}[x]\}$ and then a basis $\{p_1(M), \dots, p_s(M)\}$ for $P \cap T$. Then $M^n \in T \iff M^n \in P \cap T$. We call this the *polynomial version* of the matrix power problem. Observe that $\dim(V) \geq \dim(T) \geq \dim(T \cap P)$, so the dimension of the target vector space does not grow during the described reductions. All described operations may be performed in polynomial time using standard techniques from linear algebra.

2.2 Singular and nonsingular problem instances

An instance (A, x, V) of the orbit problem is defined as *non-singular* if no quotient of two eigenvalues of A is a root of unity. In general, it is possible to reduce an arbitrary orbit problem instance to a set of non-singular instances as follows.

Let L be the least common multiple of the orders of the quotients which are roots of unity. For each $i \in \{0, \dots, L-1\}$, we consider separately the problem of deciding whether there exists $n \in \mathbb{N}$ such that $(A^L)^n (A^i x) \in V$. The original problem instance is positive if and only if at least one of these L instances is positive.

It is easy to see that these instances are all non-singular. The eigenvalues of A^L are exactly λ_i^L where λ_i are the eigenvalues of A . If for any two distinct such eigenvalues, say $\lambda_i^L \neq \lambda_j^L$, we have $(\lambda_i^L / \lambda_j^L)^t = 1$, then λ_i / λ_j must also be a root of unity. Then by the definition of L , $\lambda_i^L / \lambda_j^L = 1$, which gives the contradiction $\lambda_i^L = \lambda_j^L$.

Observe that the reduction to the matrix power problem in the previous section preserves non-singularity, since the eigenvalues of the invertible matrix M are a subset of the eigenvalues of A . However, the non-singularity comes at the cost of higher complexity, as L may be exponentially large in the size of A . A deterministic machine would have to examine L problem instances to decide the original instance. Even if we allow non-determinism, we must still calculate A^L , whose entries require space exponential in the size of the original problem instance.

We show this reduction to make the point that if one only aims to prove decidability for the orbit problem, without regard for complexity, non-singularity of the problem instance may be assumed without loss of generality. Below we give tighter complexity upper bounds for the orbit problem with a target space of dimension at most 3 and explicitly handle singular problem instances, but if we only aimed at proving decidability, the argument could be shortened significantly.

Finally, we point out that a similar trick was used by Vereshchagin [9] to prove decidability of Skolem's problem for recurrences of depth 3 and 4. Given

a singular recurrent sequence $u(n)$, one may take

$$L = \text{lcm} \{m \mid \lambda_i/\lambda_j \text{ is an } m\text{-th root of unity}\}$$

where λ_i are the roots of the characteristic polynomial of the sequence. Then each sequence $v_i(n) = u(Ln + i)$ for $i \in \{0, \dots, L-1\}$ is non-singular in the sense that dividing two of its roots does not give a root of unity.

2.3 Towards a system of equations

Suppose now we have an instance (A, p_1, \dots, p_s) of the polynomial version of the matrix power problem. Calculate the minimal polynomial $f_A(x)$ of A and obtain canonical representations of its roots $\alpha_1, \dots, \alpha_k$, that is, the eigenvalues of A . This may be done in polynomial time, see Appendix A. Throughout this paper, for an eigenvalue α_i we will denote by $\text{mul}(\alpha_i)$ the multiplicity of α_i in the minimal polynomial of the matrix.

Fix an exponent $n \in \mathbb{N}$ and coefficients $a_1, \dots, a_s \in \mathbb{C}$ and define the polynomials $P(x) = \sum_{i=1}^s a_i p_i(x)$ and $Q(x) = x^n$. It is easy to see that

$$Q(A) = P(A)$$

if and only if

$$\forall i \in \{1, \dots, k\}. \forall j \in \{0, \dots, \text{mul}(\alpha_i) - 1\}. P^{(j)}(\alpha_i) = Q^{(j)}(\alpha_i) \quad (1)$$

Indeed, $P - Q$ is zero at A if and only if $f_A(x)$ divides $P - Q$, that is, each α_i is a root of $P - Q$ with multiplicity at least $\text{mul}(\alpha_i)$. This is equivalent to saying that each α_i is a root of $P - Q$ and its first $\text{mul}(\alpha_i) - 1$ derivatives.

Thus, in order to decide whether there exists an exponent n and coefficients a_i such that $A^n = \sum_{i=1}^s a_i p_i(A)$, it is sufficient to solve a system of equations (1) where the unknowns are n and a_1, \dots, a_s . Each eigenvalue α_i contributes $\text{mul}(\alpha_i)$ equations which specify that $P(x)$ and its first $\text{mul}(\alpha_i) - 1$ derivatives all vanish at α_i .

For brevity in what follows, we will denote by $eq(\alpha_i, j)$ the j -th derivative equation contributed to the system by α_i , that is, $P^{(j)}(\alpha_i) = Q^{(j)}(\alpha_i)$. This notation is meaningful only for $0 \leq j < \text{mul}(\alpha_i)$. We will also denote by $Eq(\alpha_i)$ the set of equations contributed by α_i to the system:

$$Eq(\alpha_i) = \{eq(\alpha_i, 0), \dots, eq(\alpha_i, \text{mul}(\alpha_i) - 1)\}$$

For example, if $f_A(x)$ has roots $\alpha_1, \alpha_2, \alpha_3$ with multiplicities $\text{mul}(\alpha_i) = i$ and the target space is $\text{span}\{p_1(A), p_2(A)\}$ then the system contains six equations:

$$\begin{aligned} \alpha_1^n &= a_1 p_1(\alpha_1) + a_2 p_2(\alpha_1) \\ \alpha_2^n &= a_1 p_1(\alpha_2) + a_2 p_2(\alpha_2) \\ n\alpha_2^{n-1} &= a_1 p_1'(\alpha_2) + a_2 p_2'(\alpha_2) \\ \alpha_3^n &= a_1 p_1(\alpha_3) + a_2 p_2(\alpha_3) \\ n\alpha_3^{n-1} &= a_1 p_1'(\alpha_3) + a_2 p_2'(\alpha_3) \\ n(n-1)\alpha_3^{n-2} &= a_1 p_1''(\alpha_3) + a_2 p_2''(\alpha_3) \end{aligned}$$

Then $Eq(\alpha_3, 0)$ is the equation

$$\alpha_3^n = a_1 p_1(\alpha_3) + a_2 p_2(\alpha_3)$$

and $Eq(\alpha_2)$ is the two equations

$$\alpha_2^n = a_1 p_1(\alpha_2) + a_2 p_2(\alpha_2)$$

$$n\alpha_2^{n-1} = a_1 p_1'(\alpha_2) + a_2 p_2'(\alpha_2)$$

3 One-dimensional version

Suppose we are given a one-dimensional matrix power problem instance (A, p) and wish to decide whether $A^n \in \text{span}\{p(A)\}$ for some n . We have constructed a system in the exponent n and the coefficient a as in (1). For example, if the roots of $f_A(x)$ are $\alpha_1, \alpha_2, \alpha_3$ with multiplicities $\text{mul}(\alpha_i) = i$, the system is:

$$\begin{aligned}\alpha_1^n &= ap(\alpha_1) \\ \alpha_2^n &= ap(\alpha_2) \\ n\alpha_2^{n-1} &= ap'(\alpha_2) \\ \alpha_3^n &= ap(\alpha_3) \\ n\alpha_3^{n-1} &= ap'(\alpha_3) \\ n(n-1)\alpha_3^{n-2} &= ap''(\alpha_3)\end{aligned}$$

In this section we will describe how such systems may be solved in polynomial time. We allow the problem instance to be singular, that is, the ratios of eigenvalues of A may be roots of unity.

The strategy is to consider quotients of equations. This eliminates the unknown coefficient a and leaves only the exponent n . First, we perform some preliminary calculations.

1. We check whether $a = 0$ has a corresponding n which solves the matrix equation $A^n = ap(A)$. This may be done using Kannan and Lipton's algorithm for the original orbit problem. If this is the case, we are done. Otherwise, assume $a \neq 0$.
2. Let $c = \max_i \{\text{mul}(\alpha_i)\}$. We check for all $n < c$ whether A^n is a multiple of $p(A)$. If so, we are done. Otherwise, assume $n \geq c$.
3. We check whether $\alpha_i = 0$ for some i . If so, then all of the equations $Eq(\alpha_i)$ are of the form $0 = ap^{(t)}(0)$, which is equivalent to $0 = p^{(t)}(0)$. We can easily check whether these equations are satisfied. If so, we dismiss them from the system without changing the set of solutions. If not, then we are done. Now we assume $\alpha_i \neq 0$ for all i .
4. Finally, we check whether the right hand side $ap^{(t)}(\alpha_i)$ of some equation is equal to 0, using a polynomial division of $p^{(t)}(x)$ by the minimal polynomial of α_i . If this is the case, then the problem instance is negative, because the left-hand sides are all non-zero.

Let $eq(\alpha_i, k) / eq(\alpha_j, t)$ denote the equation obtained by dividing $eq(\alpha_i, k)$ by $eq(\alpha_j, t)$, that is,

$$\frac{n(n-1)\dots(n-k+1)\alpha_i^{n-k}}{n(n-1)\dots(n-t+1)\alpha_j^{n-t}} = \frac{p^{(k)}(\alpha_i)}{p^{(t)}(\alpha_j)}$$

We compute canonical representations of all quotients α_i/α_j , and consider three cases.

Case 1. Some quotient α_i/α_j is not a root of unity. Then $eq(\alpha_i, 0)$ and $eq(\alpha_j, 0)$ together imply $eq(\alpha_i, 0) / eq(\alpha_j, 0)$, that is,

$$\left(\frac{\alpha_i}{\alpha_j}\right)^n = \frac{p(\alpha_i)}{p(\alpha_j)}$$

By Lemma 1 in Appendix A, we can compute canonical representations of $p(\alpha_i)/p(\alpha_j)$ and α_i/α_j in polynomial time. Then by Lemma 5 in Appendix C, n is bounded by a polynomial in the input. We check $A^n \in \text{span}\{p(A)\}$ for all n up to the bound and we are done.

Case 2. All quotients α_i/α_j are roots of unity, and all roots of $f_A(x)$ are simple. Then the system is equivalent to

$$a = \frac{\alpha_1^n}{p(\alpha_1)} \wedge \bigwedge_{i < j} \frac{eq(\alpha_i, 0)}{eq(\alpha_j, 0)}$$

It is sufficient to determine whether there exists some n which satisfies

$$\bigwedge_{i < j} \frac{eq(\alpha_i, 0)}{eq(\alpha_j, 0)} \quad (2)$$

Consider each equation $eq(\alpha_i, 0) / eq(\alpha_j, 0)$:

$$\left(\frac{\alpha_i}{\alpha_j}\right)^n = \frac{p(\alpha_i)}{p(\alpha_j)} \quad (3)$$

Suppose α_i/α_j is an r -th root of unity. If the right-hand side of (3) is also an r -th root of unity, then the solutions of (3) are $n \equiv t \pmod r$ for some t . If not, then (3) has no solution, so the entire system (1) has no solution, and the problem instance is negative. By Lemma 1, we can determine in polynomial time whether the right-hand side of (3) is a root of unity, and if so, calculate t . We transform each equation in (2) into an equivalent congruence in n . This gives a system of congruences in n which is equivalent to (2). We solve it using the Chinese Remainder Theorem. The problem instance is positive if and only if the system of congruences has a solution.

Case 3. All quotients α_i/α_j are roots of unity, and $f_A(x)$ has repeated roots. We transform the system into an equivalent one in the following way. First, we include in the new system all the quotients of equations $eq(\alpha_i, 0)$

as in Case 2. Second, for each repeated root α_i of $f_A(x)$, we take the quotients $\bigwedge_{j=0}^{mul(\alpha_i)-2} eq(\alpha_i, j) / eq(\alpha_i, j+1)$. Third, we include the equation $a = \alpha_1 / p(\alpha_1)$.

$$\bigwedge_{i < j} \frac{eq(\alpha_i, 0)}{eq(\alpha_j, 0)} \wedge \bigwedge_i \bigwedge_{j=0}^{mul(\alpha_i)-2} \frac{eq(\alpha_i, j)}{eq(\alpha_i, j+1)} \wedge a = \frac{\alpha_1}{p(\alpha_1)}$$

We solve the first part of the system as in Case 2. If there is no solution, then we are done. Otherwise, the solution is some congruence $n \equiv t_1 \pmod{t_2}$. For the remainder of the system, each ratio $eq(\alpha_i, j) / eq(\alpha_i, j+1)$ contributed by a repeated root α_i has the shape

$$\frac{\alpha_i}{n-j} = \frac{p^{(j)}(\alpha_i)}{p^{(j+1)}(\alpha_i)}$$

which is equivalent to

$$n = j + \frac{p^{(j+1)}(\alpha_i)}{p^{(j)}(\alpha_i)} \alpha_i \quad (4)$$

For each such equation (4), we calculate the right-hand side and check whether it is in \mathbb{N} . If not, then the system has no solution. Otherwise, (4) points to a single candidate n_0 . We do this for all equations where n appears outside the exponent. If they point to the same value of n , then the system is equivalent to

$$\begin{aligned} n &\equiv t_1 \pmod{t_2} \\ n &= n_0 \\ a &= \alpha_1^n / q(\alpha_1) \end{aligned}$$

We check whether n_0 satisfies the congruence and we are done.

4 Two-dimensional version

Suppose we are given (A, p_1, p_2) where A is a square matrix and p_1, p_2 are polynomials. We wish to decide whether there exists $n \in \mathbb{N}$ such that $A^n \in \text{span}\{p_1(A), p_2(A)\}$. In this section we will show that this problem is in the complexity class NP^{EqSLP} , and hence in NP^{RP} . The instance is allowed to be singular, that is, there may exist distinct eigenvalues of A whose ratio is a root of unity. We have derived a system of equations (1).

The broad strategy will be to choose a tuple of equations from the system (1), obtain a Skolem instance of depth 3 from this tuple and hence obtain a bound m on the exponent n such that if

$$A^n \in \text{span}\{p_1(A), p_2(A)\}$$

then $n < m$. This bound will be at most exponential in the size of the input, so an NP machine will be able to guess an exponent n up to the bound. Then

the machine calculates A^n , representing numbers as arithmetic circuits, and expresses membership in the target vector space as an instance of the EqSLP problem: determining whether a given arithmetic circuit evaluates to 0.

For example, suppose A has three distinct eigenvalues α, β, γ , and consider the tuple of equations $eq(\alpha, 0)$, $eq(\beta, 0)$, $eq(\gamma, 0)$:

$$\begin{pmatrix} \alpha^n \\ \beta^n \\ \gamma^n \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\beta) \\ p_1(\gamma) \end{pmatrix} + a_2 \begin{pmatrix} p_2(\alpha) \\ p_2(\beta) \\ p_2(\gamma) \end{pmatrix}$$

If the vectors on the right-hand side are linearly independent, then this triple states that the point in \mathbb{A}^3 described by the left-hand side lies on the plane in \mathbb{A}^3 described by the right-hand side. We can calculate the normal $(A_1, A_2, A_3)^T$ of the plane to obtain

$$A_1\alpha^n + A_2\beta^n + A_3\gamma^n = 0$$

Now observe that the left-hand side as a function of n satisfies a linear recurrence of depth 3 over \mathbb{A} . Provided that the ratios of α, β, γ are not roots of unity, Lemmas 6, 7, 8 give a bound on n which is at most exponential in the input size, as desired. Problems arise, however, when this linear recurrence is allowed to be singular. For example, suppose that $A_3 = 0$, and let α/β and $-A_2/A_1$ be roots of unity of the same order. Then the zeroes of the recurrence are a full arithmetic progression, and this Skolem instance fails to give a bound on n .

Similarly, if the matrix A has only two eigenvalues α, β , but one of them, say α , is repeated, then we may consider the triple $eq(\alpha, 0)$, $eq(\alpha, 1)$, $eq(\beta, 0)$. Using similar reasoning, we see that

$$A_1\alpha^n + A_2n\alpha^{n-1} + A_3\beta^n = 0$$

must hold for some effective algebraic constants A_1, A_2, A_3 . This corresponds to a Skolem instance of depth 3 where one of the characteristic roots is repeated. Now provided that α/β is not a root of unity, we have an exponential bound on n from Lemma 9. However, if α/β and $-A_3/A_1$ are roots of unity, and $A_2 = 0$, then the Skolem instance could have infinitely many solutions, failing to give a bound on n .

In order to explicitly handle the case of singular orbit instances, we will consider the relation \sim on the eigenvalues of A , defined by

$$\alpha \sim \beta \text{ if and only if } \alpha/\beta \text{ is a root of unity}$$

It is clear that \sim is an equivalence relation. The equivalence classes $\mathcal{C}_1, \dots, \mathcal{C}_k$ of \sim are of two kinds. First, a class can be its own image under complex conjugation:

$$\mathcal{C}_i = \{\bar{\alpha} \mid \alpha \in \mathcal{C}_i\}$$

Each such self-conjugate class $\{\alpha_1, \dots, \alpha_s\}$ has the form $\{\alpha\omega_1, \dots, \alpha\omega_s\}$ where ω_i are roots of unity, and $|\alpha_j| = \alpha \in \mathbb{R} \cap \mathbb{A}$. Call this α the *stem* of the equivalence class \mathcal{C}_i . Second, if an equivalence class is not self-conjugate, then its image

under complex conjugation must be another equivalence class of \sim . Thus, the remaining equivalence classes of \sim are grouped into pairs $(\mathcal{C}_i, \mathcal{C}_j)$ such that $\mathcal{C}_i = \{\bar{x} \mid x \in \mathcal{C}_j\} = \overline{\mathcal{C}_j}$. In this case, we can write \mathcal{C}_i and \mathcal{C}_j as

$$\begin{aligned}\mathcal{C}_i &= \{\lambda\omega_1, \dots, \lambda\omega_s\} \\ \mathcal{C}_j &= \{\overline{\lambda\omega_1}, \dots, \overline{\lambda\omega_s}\}\end{aligned}$$

where ω_i are roots of unity and $\lambda \in \mathbb{A}$ with $\arg(\lambda)/2\pi \notin \mathbb{Q}$. Call λ the *stem* of \mathcal{C}_i and $\bar{\lambda}$ the stem of \mathcal{C}_j . Observe that the stems of self-conjugate classes are distinct positive real numbers, and that the ratio $\lambda/\bar{\lambda}$ of the stems of paired classes cannot be a root of unity. Recall also that we can assume the eigenvalues of A are algebraic integers, as a by-product of the reduction from the orbit problem. Since roots of unity and their inverses are algebraic integers, it follows that the stems of equivalence classes must also be algebraic integers.

Let

$$Eq(\mathcal{C}) = \bigcup_{\alpha \in \mathcal{C}} Eq(\alpha)$$

denote the set of equations contributed to the system by the eigenvalues in \mathcal{C} , and let

$$Eq(\mathcal{C}, i) = \bigcup_{\substack{\alpha \in \mathcal{C} \\ i < mul(\alpha)}} \{eq(\alpha, i)\}$$

denote the set of i -th derivative equations contributed by the roots in \mathcal{C} .

We will case split on the equivalence classes of \sim .

Case I. Suppose \sim has exactly one equivalence class $\mathcal{C} = \{\alpha\omega_1, \dots, \alpha\omega_s\}$, necessarily self-conjugate, with stem α . Let L be the least common multiple of the orders of $\omega_1, \dots, \omega_s$. Since L is at most exponentially large in the size of the input, it can be written using polynomially bits, so our NP machine will guess $r = n \bmod L$. Now consider the set of equations $Eq(\mathcal{C}, 0)$:

$$\begin{aligned}(\alpha\omega_1)^n &= a_1p(\alpha\omega_1) + a_2p(\alpha\omega_1) \\ &\vdots \\ (\alpha\omega_s)^n &= a_1p(\alpha\omega_s) + a_2p(\alpha\omega_s)\end{aligned}$$

Having guessed r , we can easily calculate $\omega_1^n, \dots, \omega_s^n$ in polynomial time, since ω_i are roots of unity whose order divides L . Then the equations $Eq(\mathcal{C}, 0)$ are equivalent to

$$\begin{pmatrix} \alpha^n \\ \vdots \\ \alpha^n \end{pmatrix} = B \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \tag{5}$$

where B is an $s \times 2$ matrix over \mathbb{A} , computable in polynomial time. Next we subtract the first row of (5) from rows $2, \dots, s$, obtaining

$$\alpha^n = \varphi_1 a_1 + \varphi_2 a_2 \wedge \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix} = B' \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}$$

Here $(\varphi_1 \ \varphi_2)$ is the first row of the matrix B , and B' is the bottom $s - 1$ rows of B . Thus, $Eq(\mathcal{C}, 0)$ is equivalent to

$$\alpha^n = \varphi_1 a_1 + \varphi_2 a_2$$

together with the constraint that $(a_1 \ a_2)^T$ must lie in the nullspace of B' . We calculate the nullspace of B' directly. If its dimension is less than 2, then we have a linear constraint on a_1, a_2 . This constraint is of the form $a_1 = ka_2$ when the nullspace of B' has dimension 1, and is $a_1 = a_2 = 0$ when the nullspace is of dimension 0. In both cases, we substitute into the system (1), and solve the resulting lower-dimensional system using the algorithms for the one-dimensional orbit problem and Kannan and Lipton's original orbit problem. In the case when the nullspace of B' has dimension 2, then the linear constraint is vacuous, and $Eq(\mathcal{C}, 0)$ is equivalent to $\alpha^n = \varphi_1 a_1 + \varphi_2 a_2$.

In the same way we collapse $Eq(\mathcal{C}, 1)$ into one first-derivative equation

$$n\alpha^{n-1} = \varphi_3 a_1 + \varphi_4 a_2$$

using our guess for $n \bmod L$. We do this for all $Eq(\mathcal{C}, i)$, obtaining a system of equations equivalent to (1) based on the stem of \mathcal{C} , rather than the actual eigenvalues in \mathcal{C} . Denote the resulting set of equations by $\mathcal{F}(Eq(\mathcal{C}))$.

If some eigenvalue $x \in \mathcal{C}$ has $\text{mul}(x) \geq 3$, then $\mathcal{F}(Eq(\mathcal{C}))$ contains the following triple of equations:

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \\ n(n-1)\alpha^{n-2} \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \\ \varphi_5 \end{pmatrix} + a_2 \begin{pmatrix} \varphi_2 \\ \varphi_4 \\ \varphi_6 \end{pmatrix} \quad (6)$$

If the vectors on the right-hand side of (6) are linearly independent, then they specify a plane in \mathbb{A}^3 , and the triple states that the point on the left-hand side must lie on this plane. We calculate the normal $(A_1 \ A_2 \ A_3)^T$ of the plane and obtain

$$A_1 \alpha^n + A_2 n \alpha^{n-1} + A_3 n(n-1) \alpha^{n-2} = 0$$

This is a quadratic equation in n . It has at most two solutions, both at most exponentially large in the size of the input, so we are done. If the vectors on the right-hand side of (6) are linearly dependent, then we may equivalently consider

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \\ n(n-1)\alpha^{n-2} \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \\ \varphi_5 \end{pmatrix}$$

We divide the first equation by the second to obtain

$$\frac{\alpha}{n} = \frac{\varphi_1}{\varphi_3}$$

which limits n at most one, exponentially large, candidate value.

If all eigenvalues x in \mathcal{C} have $\text{mul}(x) \leq 2$ and at least one has $\text{mul}(x) = 2$, then $\mathcal{F}(Eq(\mathcal{C}))$ consists of exactly two equations:

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \end{pmatrix} + a_2 \begin{pmatrix} \varphi_2 \\ \varphi_4 \end{pmatrix} \quad (7)$$

If $(\varphi_1 \ \varphi_3)^T$ and $(\varphi_2 \ \varphi_4)^T$ are linearly independent, then the right-hand side of (7) spans all of \mathbb{A}^2 as a_1, a_2 range over \mathbb{A} , so the problem instance is trivially positive. Otherwise, we may equivalently consider

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \end{pmatrix}$$

and limit n to at most one candidate value which is exponentially large in the input size.

Finally, if all eigenvalues x in \mathcal{C} have $\text{mul}(x) = 1$, then $\mathcal{F}(Eq(\mathcal{C}))$ contains only the equation

$$\alpha^n = a_1\varphi_1 + a_2\varphi_2$$

which has a solution if and only if at least one of φ_1, φ_2 is non-zero.

Case II. Suppose \sim has exactly two equivalence classes, \mathcal{C}_1 and \mathcal{C}_2 , with respective stems α and β , so that

$$\mathcal{C}_1 = \{\alpha\omega_1, \dots, \alpha\omega_s\}$$

$$\mathcal{C}_2 = \{\beta\omega'_1, \dots, \beta\omega'_l\}$$

The classes could be self-conjugate, in which case $\alpha, \beta \in \mathbb{A} \cap \mathbb{R}$, or they could be each other's image under complex conjugation, in which case $\alpha = \overline{\beta}$. In both cases, α/β is not a root of unity.

As in the previous case, we define L to be the least common multiple of the orders of $\omega_1, \dots, \omega_s, \omega'_1, \dots, \omega'_l$, and our NP machine guesses $r = n \bmod L$. We transform the system $Eq(\mathcal{C}_1) \wedge Eq(\mathcal{C}_2)$ into the equivalent system $\mathcal{F}(Eq(\mathcal{C}_1)) \wedge \mathcal{F}(Eq(\mathcal{C}_2))$. If all eigenvalues x of A have $\text{mul}(x) = 1$, then the system consists of two equations, one for each equivalence class of \sim :

$$\begin{aligned} \alpha^n &= a_1\varphi_1 + a_2\varphi_2 \\ \beta^n &= a_1\varphi_3 + a_2\varphi_4 \end{aligned}$$

If $(\varphi_1 \ \varphi_3)^T$ and $(\varphi_2 \ \varphi_4)^T$ are linearly independent, then the problem instance is trivially positive. Otherwise, it suffices to look for n which satisfies

$$\begin{aligned} \alpha^n &= a_1\varphi_1 \\ \beta^n &= a_1\varphi_3 \end{aligned}$$

and hence

$$\left(\frac{\alpha}{\beta}\right)^n = \frac{\varphi_1}{\varphi_3}$$

A bound on n follows from Lemma 5. This argument relies crucially on the fact that α/β is not a root of unity.

If some eigenvalue x of A has $\text{mul}(x) \geq 2$, say $x \in \mathcal{C}_1$, then the system contains the following triple of equations:

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \\ \beta^n \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \\ \varphi_5 \end{pmatrix} + a_2 \begin{pmatrix} \varphi_2 \\ \varphi_4 \\ \varphi_6 \end{pmatrix} \quad (8)$$

If the vectors on the right-hand side of (8) are linearly dependent, so that the right-hand side describes a space of dimension 1, it suffices to look for solutions to

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \\ \beta^n \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \\ \varphi_5 \end{pmatrix}$$

Then dividing we obtain

$$\frac{\alpha}{n} = \frac{\varphi_1}{\varphi_3}$$

which limits n to at most one, exponentially large candidate value. Otherwise, if the vectors on the right-hand side of (8) are linearly independent, we calculate the normal $(A_1 \ A_2 \ A_3)^T$ to the plane described by them and obtain

$$A_1\alpha^n + A_2n\alpha^{n-1} + A_3\beta^n = 0$$

A bound on n which is exponential in the size of the input follows from Lemma 9. This relies on the fact that α/β cannot be a root of unity.

Case III. Suppose \sim has at least three equivalence classes. Then we can choose eigenvalues α, β, γ , each from a distinct equivalence class. Then consider $eq(\alpha, 0)$, $eq(\beta, 0)$ and $eq(\gamma, 0)$:

$$\begin{pmatrix} \alpha^n \\ \beta^n \\ \gamma^n \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\beta) \\ p_1(\gamma) \end{pmatrix} + a_2 \begin{pmatrix} p_2(\alpha) \\ p_2(\beta) \\ p_2(\gamma) \end{pmatrix}$$

If the vectors on the right-hand side are linearly independent, we calculate the normal $(A_1, A_2, A_3)^T$ of the plane on the right-hand side to obtain

$$A_1\alpha^n + A_2\beta^n + A_3\gamma^n = 0$$

This is a non-singular Skolem instance of depth 3, so a bound on n follows from Lemmas 6, 7, 8. If the vectors on the right-hand side are not linearly independent, then we may equivalently consider

$$\begin{pmatrix} \alpha^n \\ \beta^n \\ \gamma^n \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\beta) \\ p_1(\gamma) \end{pmatrix}$$

which gives

$$\left(\frac{\alpha}{\beta}\right)^n = \frac{p_1(\alpha)}{p_1(\beta)}$$

An exponential bound on n follows from Lemma 5, because α/β is not a root of unity.

Thus, we have now shown that in all cases, an NP machine may compute a bound m such that if

$$A^n x \in \text{span}\{y, z\}$$

then $n < m$. This bound is at most exponential in the size of the input. From here it is easy to argue membership in NP^{EqSLP} . The machine guesses some n up to the bound, ensuring that this n is consistent with the guess for $n \bmod L$. Now we need to compute $A^n x$ and check whether it is in the target vector space. Since n is exponential in the size of the input, the entries of $A^n x$ are, in general, doubly-exponential. That is, they require an exponential number of bits to write down. However, the entries of $A^n x$ may easily be represented as polynomial-sized arithmetic circuits. We consider all projections of $A^n x$, y and z to three coordinates, and for each projection we expression the question of linear independence as the zeroness of a 3×3 determinant, also expressed as an arithmetic circuit. It is clear that n is a witness to the problem instance if and only if for any projection to three coordinates, $A^n x$, y and z are linearly dependent. This is easy to determine with an EqSLP oracle, so we have membership in NP^{EqSLP} . It is known that $\text{EqSLP} \subseteq \text{coRP}$ [14], so we also have membership in NP^{RP} .

5 Three-dimensional version

Suppose we have a problem instance (A, p_1, p_2, p_3) and wish to decide whether $A^n \in \text{span}\{p_1(A), p_2(A), p_3(A)\}$ for some n . As before, we have constructed a system (1) in n and the coefficients a_1, a_2, a_3 . The eigenvalues of A are algebraic integers and do not include 0. In this section we will show that there exists an effective bound m which is at most exponentially large in the size of the input, such that

$$A^n \in \text{span}\{p_1(A), p_2(A), p_3(A)\} \Rightarrow n < m$$

Then by the same reasoning as in the two-dimensional case, we will have membership in NP^{EqSLP} and NP^{RP} for the three-dimensional orbit problem.

Following the strategy of the two-dimensional case, we will select tuples of equations and obtain a bound on n using the lemmas for Skolem's problem for recurrences of depth 4 in Appendix F. We will again case split on the equivalence classes of the relation \sim .

Case I. Suppose there are at least two pairs of classes $(\mathcal{C}_i, \overline{\mathcal{C}_i})$, $(\mathcal{C}_j, \overline{\mathcal{C}_j})$ which are not self-conjugate. Then let $\alpha \in \mathcal{C}_i$, $\beta = \overline{\alpha} \in \overline{\mathcal{C}_i}$, $\gamma \in \mathcal{C}_j$, $\delta = \overline{\gamma} \in \overline{\mathcal{C}_j}$. Then we

consider the tuple of equations

$$\begin{pmatrix} \alpha^n \\ \beta^n \\ \gamma^n \\ \delta^n \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\beta) \\ p_1(\gamma) \\ p_1(\delta) \end{pmatrix} + a_2 \begin{pmatrix} p_2(\alpha) \\ p_2(\beta) \\ p_2(\gamma) \\ p_2(\delta) \end{pmatrix} + a_3 \begin{pmatrix} p_3(\alpha) \\ p_3(\beta) \\ p_3(\gamma) \\ p_3(\delta) \end{pmatrix} \quad (9)$$

If the vectors on the right-hand side are linearly dependent, then we rewrite the right-hand side as a linear combination of $h < 3$ vectors and obtain a bound on n as we did for tuples of equations in the one- and two-dimensional orbit problem. If the vectors on the right-hand side of (9) are linearly independent, then we calculate the normal of their three-dimensional subspace of \mathbb{A}^4 , obtaining an equation

$$A_1\alpha^n + A_2\beta^n + A_3\gamma^n + A_4\delta^n = 0 \quad (10)$$

and hence an exponential bound on n from Lemmas 12 and 13. We are relying on the fact that the ratios of $\alpha, \beta, \gamma, \delta$ are not roots of unity. Notice that we need (α, β) and (γ, δ) to be pairwise complex conjugates in order to apply Lemma 13.

Case II. Suppose now that there is exactly one pair of classes $(\mathcal{C}_i, \overline{\mathcal{C}}_i)$ which are not self-conjugate. In general, for any eigenvalue x of A we must have $\text{mul}(x) = \text{mul}(\overline{x})$. Therefore, if any eigenvalue $\alpha \in \mathcal{C}_i$ has $\text{mul}(\alpha) > 1$, we can select the non-singular tuple of equations $\text{eq}(\alpha, 0)$, $\text{eq}(\alpha, 1)$, $\text{eq}(\overline{\alpha}, 0)$, $\text{eq}(\overline{\alpha}, 1)$:

$$\begin{pmatrix} \alpha^n \\ \overline{\alpha}^n \\ n\alpha^{n-1} \\ n\overline{\alpha}^{n-1} \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\overline{\alpha}) \\ p'_1(\alpha) \\ p'_1(\overline{\alpha}) \end{pmatrix} + a_2 \begin{pmatrix} p_2(\alpha) \\ p_2(\overline{\alpha}) \\ p'_2(\alpha) \\ p'_2(\overline{\alpha}) \end{pmatrix} + a_3 \begin{pmatrix} p_3(\alpha) \\ p_3(\overline{\alpha}) \\ p'_3(\alpha) \\ p'_3(\overline{\alpha}) \end{pmatrix}$$

This leads to a non-singular Skolem instance of depth 4 over \mathbb{A} for a recurrence sequence with two repeated characteristic roots:

$$A_1\alpha^n + A_2\overline{\alpha}^n + A_3n\alpha^{n-1} + A_4n\overline{\alpha}^{n-1} = 0$$

An exponential bound on n follows from Lemma 10, since $\alpha/\overline{\alpha}$ is not a root of unity.

We can now assume that eigenvalues in \mathcal{C}_i and $\overline{\mathcal{C}}_i$ contribute exactly one equation to the system. Now we perform the collapsing operation used in the two-dimensional orbit problem, transforming $\text{Eq}(\mathcal{C}_i) \wedge \text{Eq}(\overline{\mathcal{C}}_i)$ into the equivalent $\mathcal{F}(\text{Eq}(\mathcal{C}_i)) \wedge \mathcal{F}(\text{Eq}(\overline{\mathcal{C}}_i))$. Since all eigenvalues in \mathcal{C}_i and $\overline{\mathcal{C}}_i$ contribute one equation each, $\mathcal{F}(\text{Eq}(\mathcal{C}_i)) \wedge \mathcal{F}(\text{Eq}(\overline{\mathcal{C}}_i))$ is just

$$\begin{aligned} \lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3 \\ \overline{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6 \end{aligned}$$

where $\lambda, \overline{\lambda}$ are the stems of \mathcal{C}_i and $\overline{\mathcal{C}}_i$. We perform the collapsing operation to all self-conjugate classes as well, reducing the system of equations to an equivalent system based on the stems of the equivalence classes, not the actual eigenvalues of A . This is beneficial, because the stems cannot divide to give roots of unity,

so we can use 4-tuples of equations to construct non-singular Skolem instances of order 4.

If there are at least two self-conjugate equivalence classes, with respective stems α, β , we take the tuple

$$\begin{aligned}\lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3 \\ \bar{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6 \\ \alpha^n &= a_1\varphi_7 + a_2\varphi_8 + a_3\varphi_9 \\ \beta^n &= a_1\varphi_{10} + a_2\varphi_{11} + a_3\varphi_{12}\end{aligned}$$

and obtain the non-singular Skolem instance

$$A_1\lambda^n + A_2\bar{\lambda}^n + A_3\alpha^n + A_4\beta^n = 0$$

Then we have a bound on n from Lemmas 12 and 13. Similarly, if there is only one self-conjugate equivalence class, with stem α , but some of its eigenvalues are repeated, we use the tuple

$$\begin{aligned}\lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3 \\ \bar{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6 \\ \alpha^n &= a_1\varphi_7 + a_2\varphi_8 + a_3\varphi_9 \\ n\alpha^{n-1} &= a_1\varphi_{10} + a_2\varphi_{11} + a_3\varphi_{12}\end{aligned}$$

to obtain the non-singular instance

$$A_1\lambda^n + A_2\bar{\lambda}^n + A_3\alpha^n + A_4n\alpha^{n-1} = 0$$

which gives a bound on n according to Lemma 11. If there is exactly one self-conjugate class, with stem α , containing no repeated roots, then the system consists of three equations:

$$\begin{aligned}\lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3 \\ \bar{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6 \\ \alpha^n &= a_1\varphi_7 + a_2\varphi_8 + a_3\varphi_9\end{aligned}$$

Depending on whether the vectors $(\varphi_1 \ \varphi_4 \ \varphi_7)^T$, $(\varphi_2 \ \varphi_5 \ \varphi_8)^T$, $(\varphi_3 \ \varphi_6 \ \varphi_9)^T$ are linearly independent, this is either a trivially positive instance, or a lower-dimensional non-singular instance. Finally, if there are no self-conjugate classes, the system consists of only two equations:

$$\begin{aligned}\lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3 \\ \bar{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6\end{aligned}$$

Again, depending on the dimension of

$$\text{span} \left\{ \begin{pmatrix} \varphi_1 \\ \varphi_4 \end{pmatrix}, \begin{pmatrix} \varphi_2 \\ \varphi_5 \end{pmatrix}, \begin{pmatrix} \varphi_3 \\ \varphi_6 \end{pmatrix} \right\}$$

this is either a trivially positive instance, or a lower-dimensional non-singular one.

Case III. All equivalence classes of \sim are self-conjugate. As above, we use the collapsing operation, and pick out 4-tuples of equations. We rely on the fact that stems of classes are distinct real algebraic numbers to ensure that the resulting Skolem instances are non-singular and give us the desired bound on n .

6 Complexity and conclusion

We have shown that the higher-dimensional orbit problem is decidable in polynomial time when $\dim(V) = 1$. We have also shown membership in $\mathbf{NP}^{\text{EqSLP}}$ in the cases $\dim(V) \in \{2, 3\}$. It is known [14] that $\text{EqSLP} \subseteq \text{coRP}$, so membership in \mathbf{NP}^{RP} follows immediately.

The connection between Skolem's problem and the orbit problem in fixed dimension is evident. A Skolem problem instance requires us to find the zeroes of a quasi-polynomial $\sum_{i=0}^t \alpha_i^n p_i(n)$, whereas an orbit problem instance is equivalent to a system of equations from which we may extract quasi-polynomial equations.

In general, given a Skolem instance

$$\exists n. x^T A^n y = 0$$

where A has size $k \times k$, we can reduce to the equivalent orbit instance

$$A^n y \in (\text{span}\{x\})^\perp$$

which has a target vector space of dimension $k - 1$. In the other direction, starting with an orbit instance with a k -dimensional target, we can construct an equivalent system of equations, extract $(k + 1)$ -tuples, obtain a normal vector of a k -dimensional subspace and hence an equation which corresponds to finding the zeroes of a recurrence of depth $k + 1$.

It is interesting to note that at first glance, the orbit problem in fixed dimension appears more difficult than Skolem's problem for a fixed depth, due to its unbounded matrix. On the contrary, the presense of more eigenvalues allows us more freedom when choosing equations from which to extract a bound on the exponent n . Repetitions in the roots of f_A simplify matters greatly by bringing n into the base position. A Skolem problem instance offers no such freedom.

Because of the similarities in the mathematics of these two problems, we believe they should be explored in conjunction. Given that Skolem's problem for depth 3 and 4 was only solved with the advent of Baker's theorem, it is unsurprising that we required the same results to prove decidability for the two- and three-dimensional orbit problem.

A Algebraic numbers and operations on them

A complex number α is *algebraic* if there exists a polynomial $p \in \mathbb{Q}[x]$ such that $p(\alpha) = 0$. The set of algebraic numbers, denoted by \mathbb{A} , is a field under the usual arithmetic operations. The *minimal polynomial* of α is the unique

monic polynomial of least degree which vanishes at α and is denoted by $f_\alpha(x)$. The *degree* of $\alpha \in \mathbb{A}$ is defined as the degree of its minimal polynomial and is denoted by n_α . The *height* of α is defined as the maximum absolute value of a coefficient in the minimal polynomial of α and is denoted by H_α . The roots of $f_\alpha(x)$ (including α) are called the *Galois conjugates* of α . The *absolute norm* of α , denoted $\mathcal{N}_{abs}(\alpha)$, is the product of the Galois conjugates of α . By Viète's laws, we have

$$\mathcal{N}_{abs}(\alpha) = (-1)^{n_\alpha} \frac{a}{b}$$

where a, b are respectively the free term and the leading coefficient of $f_\alpha(x)$. It follows that $\mathcal{N}_{abs}(\alpha) \in \mathbb{Q}$. An *algebraic integer* is an algebraic number α such that $f_\alpha \in \mathbb{Z}[x]$. The set of algebraic integers, denoted $\mathbb{A}_\mathbb{Z}$, is a ring under the usual addition and multiplication.

The *canonical representation* of an algebraic number α is its minimal polynomial $f_\alpha(x)$, along with a numerical approximation of $Re(\alpha)$ and $Im(\alpha)$ of sufficient precision to distinguish α from its Galois conjugates. More precisely, we represent α by the tuple

$$(f_\alpha, x, y, R) \in (\mathbb{Q}[x] \times \mathbb{Q} \times \mathbb{Q} \times \mathbb{Q})$$

meaning that α is the unique root of f_α inside the circle centred at (x, y) in the complex plane with radius R . A bound due to Mignotte [15] states that for roots $\alpha_i \neq \alpha_j$ of a polynomial $p(x)$,

$$|\alpha_i - \alpha_j| > \frac{\sqrt{6}}{n^{(n+1)/2} H^{n-1}} \quad (11)$$

where n and H are the degree and height of p , respectively. Thus, if R is restricted to be less than a quarter of the root separation bound, the representation is well-defined and allows for equality checking. Observe that given f_α , the remaining data necessary to describe α is polynomial in the length of the input. It is known how to obtain polynomially many bits of the roots of any $p \in \mathbb{Q}[x]$ in polynomial time [16].

From here onwards, when we say an algebraic number α is given, we will assume we have a canonical description of α . We will denote by $\|\alpha\|$ the length of this description, assuming that integers are expressed in binary and rationals are expressed as pairs of integers. Observe that $|\alpha|$ is an exponentially large quantity in $\|\alpha\|$ whereas $\ln |\alpha|$ is polynomially large. Notice also that $1/\ln |\alpha|$ is at most exponentially large in $\|\alpha\|$. For a rational a , $\|a\|$ is just the sum of the lengths of its numerator and denominator written in binary. For a polynomial $p \in \mathbb{Q}[x]$, $\|p\|$ will denote $\sum_{i=0}^n \|p_i\|$ where n is the degree of the polynomial and p_i are its coefficients.

Lemma 1. *Given canonical representations of $\alpha, \beta \in \mathbb{A}$ and a polynomial $p \in \mathbb{Q}[x]$, it is possible to compute canonical descriptions of $\alpha \pm \beta$, $\alpha\beta^{\pm 1}$ and $p(\alpha)$ in time polynomial in the length of the input (that is, in $\|\alpha\| + \|\beta\| + \|p\|$).*

Proof. The resultant of $f_\alpha(x - y)$ and $f_\beta(y)$, interpreted as polynomials in y with coefficients in $\mathbb{Q}[x]$, is a polynomial in x which vanishes at $\alpha + \beta$. We compute it in polynomial time using the Sub-Resultant algorithm (see Algorithm 3.3.7 in [17]) and factor it into irreducibles using the LLL algorithm [18]. Finally, we approximate the roots of each irreducible factor to identify the minimal polynomial of $\alpha + \beta$. The degree of $\alpha + \beta$ is at most $n_\alpha n_\beta$, while its height is bounded by $H_{\alpha+\beta} \leq H_\alpha^{n_\beta} H_\beta^{n_\alpha}$ [19]. Therefore, by (11), a polynomial number of bits suffices to describe $\alpha + \beta$ unambiguously. Similarly, we can compute canonical representations of $\alpha - \beta$, $\alpha\beta$ and α/β in polynomial time using resultants, see [19].

To calculate $p(\alpha)$ we repeatedly use addition and multiplication. It suffices to prove that all intermediate results may be represented in polynomial space. It is clear that their degrees are at most n_α , but it is not obvious how quickly the coefficients of their minimal polynomials grow. However, there is a simple reason why their representation is polynomially bounded. Let A be the companion matrix of f_α . Then $p(\alpha)$ is an eigenvalue of $p(A)$. We can calculate $p(A)$ using only polynomial space. Then from the formula

$$\det(\lambda I - p(A)) = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \prod_{i=1}^n (\lambda I - p(A))_{i, \sigma(i)}$$

it is evident that the coefficients of the characteristic polynomial of $p(A)$ are exponentially large in the length of the input, so their representation requires only polynomial space. This characteristic polynomial may be factored into irreducibles in polynomial time, so the description of $p(\alpha)$ and of all intermediate results is polynomially bounded. \square

It is trivial to check whether $\alpha = \beta$ and whether α belongs to one of $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$. It takes only polynomial time to determine whether α is a root of unity, and if so, to calculate its order and phase.

B Number fields and ideals

In this section, we define some terminology from algebraic number theory and mention textbook results. For more details, see [17, 20]. We also define the ideal-counting function v_P , which is a notion of magnitude of algebraic numbers distinct from the usual absolute value. We follow the excellent presentation of [12].

An *algebraic number field* is a field extension \mathbb{K} of \mathbb{Q} which, considered as a \mathbb{Q} -vector space, has finite dimension. This dimension is called the *degree* of the number field and is denoted by $[\mathbb{K} : \mathbb{Q}]$. The primitive element theorem states that for any number field \mathbb{K} , there exists an element $\theta \in \mathbb{K}$ such that $\mathbb{K} = \mathbb{Q}(\theta)$. Such a θ is called a *primitive element* of \mathbb{K} and satisfies $n_\theta = [\mathbb{K} : \mathbb{Q}]$. The proof of the primitive element theorem is constructive and shows how to obtain a primitive element for $\mathbb{K} = \mathbb{Q}(\alpha_1, \dots, \alpha_k)$ given $\alpha_1, \dots, \alpha_k$. There exist exactly n_θ monomorphisms from \mathbb{K} into \mathbb{C} , given by $\theta \rightarrow \theta_i$, where θ_i are the

Galois conjugates of θ . If $\alpha \in \mathbb{K}$, then $n_\alpha | n_\theta$. Moreover, if $\sigma_1, \dots, \sigma_{n_\theta}$ are the monomorphisms from \mathbb{K} into \mathbb{C} then $\sigma_1(\alpha), \dots, \sigma_{n_\theta}(\alpha)$ are exactly the Galois conjugates of α , each repeated n_θ/n_α times. The *norm of α relative to \mathbb{K}* is defined as

$$\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha) = \prod_{i=1}^{n_\theta} \sigma_i(\alpha) = (\mathcal{N}_{abs}(\alpha))^{n_\theta/n_\alpha}$$

For a number field \mathbb{K} , the set $\mathbb{K}_{\mathbb{Z}} = \mathbb{A}_{\mathbb{Z}} \cap \mathbb{K}$ forms a ring under the usual addition and multiplication. The ideals of $\mathbb{K}_{\mathbb{Z}}$ are finitely generated, and form a commutative ring under the operations

$$IJ = \{xy \mid x \in I, y \in J\}$$

$$I + J = \{x + y \mid x \in I, y \in J\}$$

with unit $[1] = \mathbb{K}_{\mathbb{Z}}$ and zero $[0] = \{0\}$. An ideal P is *prime* if $P = AB$ implies $A = P$ or $A = [1]$. A well-known fact is the fundamental theorem of ideal theory: each non-zero ideal may be represented uniquely (up to reordering) as a product of prime ideals.

This theorem gives rise to the following *ideal-counting function* $v_P : \mathbb{K}_{\mathbb{Z}} \setminus \{0\} \rightarrow \mathbb{N}$. For a fixed prime ideal P , we define $v_P(\alpha)$ to be the number of times P appears in the factorisation into prime ideals of $[\alpha]$. That is,

$$v_P(\alpha) = k \text{ if and only if } P^k \mid [\alpha] \text{ and } P^{k+1} \nmid [\alpha]$$

We also define $v_P(0) = \infty$. The function satisfies the following properties:

- $v_P(\alpha\beta) = v_P(\alpha) + v_P(\beta)$
- $v_P(\alpha + \beta) \geq \min\{v_P(\alpha), v_P(\beta)\}$
- If $v_P(\alpha) \neq v_P(\beta)$, then $v_P(\alpha + \beta) = \min\{v_P(\alpha), v_P(\beta)\}$.

For any $\alpha \in \mathbb{K}$ such that $\alpha \notin \mathbb{K}_{\mathbb{Z}}$, we can find $\beta \in \mathbb{K}_{\mathbb{Z}}$ and $n \in \mathbb{Z} \subseteq \mathbb{K}_{\mathbb{Z}}$ such that $\alpha = \beta/n$. We extend v_P to \mathbb{K} by defining $v_P(\alpha) = v_P(\beta) - v_P(n)$. The first of the three properties of v_P above guarantees that this value is independent of the choice of β, n , making the extension of v_P to \mathbb{K} well-defined. Note the extension preserves the three properties.

For an ideal $I \neq [0]$, the quotient ring $\mathbb{K}_{\mathbb{Z}}/I$ is finite. The *norm* of I , denoted $\mathcal{N}(I)$, is defined as $|\mathbb{K}_{\mathbb{Z}}/I|$. We define also $\mathcal{N}([0]) = \infty$. Notice that $\mathcal{N}(I) = 1$ if and only if $I = [1]$, otherwise $\mathcal{N}(I) \geq 2$. Each prime ideal P contains a unique prime number p , and $\mathcal{N}(P) = p^f$ for some natural number $f \geq 1$. In general,

$$|\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha)| = \mathcal{N}([\alpha]) \geq 2^{v_P(\alpha)}$$

Noting that $\mathcal{N}(P) \geq 2$ for any prime ideal P , we have

$$v_P(\alpha) \leq \log_2 |\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha)| \leq \log_2 |\mathcal{N}_{abs}(\alpha)|^d$$

where $d = [\mathbb{K} : \mathbb{Q}]$. Thus, if we are given $\mathbb{K} = \mathbb{Q}(\alpha_1, \dots, \alpha_k)$ for canonically represented algebraic numbers α_i and a canonically represented $\alpha \in \mathbb{K}$, we can

observe that d is at most polynomially large in the length of the input and $|\mathcal{N}_{abs}(\alpha)|$ is at most exponentially large in the length of the input. Therefore, $v_P(\alpha)$ is only polynomially large.

The following lemma is simple, but occurs frequently in what follows, so we state it explicitly here.

Lemma 2. *Let \mathbb{K} be a number field and $\alpha \in \mathbb{K}$ with $\alpha \notin \mathbb{K}_{\mathbb{Z}}$. Then there exists a prime ideal P of $\mathbb{K}_{\mathbb{Z}}$ such that $v_P(\alpha) \neq 0$.*

Proof. There exist $\beta \in \mathbb{K}_{\mathbb{Z}}$ and $m \in \mathbb{Z}$ such that $\alpha = \beta/m$. If $[\beta] = [m]$, then β and m are associates, so α must be a unit of $\mathbb{K}_{\mathbb{Z}}$. Since $\alpha \notin \mathbb{K}_{\mathbb{Z}}$, it follows that $[\beta] \neq [m]$, so the factorisations of $[\beta]$ and $[m]$ into prime ideals must differ. Therefore, $v_P(\beta) \neq v_P(m)$ for some prime ideal P , so $v_P(\alpha) \neq 0$. \square

C Baker's theorem and van der Poorten's theorem

Theorem 1. *(Baker [11]) Let $\alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n$ be algebraic numbers of degree at most d . Let $A_i \geq 4$ be an upper bound for the height of α_i , and $H \geq 4$ be an upper bound for the heights of β_1, \dots, β_n . If*

$$\Lambda = \beta_1 \ln \alpha_1 + \dots + \beta_n \ln \alpha_n \neq 0$$

and

$$\Omega = \ln A_1 \dots \ln A_n$$

then

$$\ln |\Lambda| > -c \ln(H\Omega) \Omega$$

where $c = (16nd)^{200n}$.

Lemma 3. *Let $\lambda, b \in \mathbb{C}$, where $|\lambda| = 1$ and λ is not a root of unity. Suppose $\phi(n)$ is a function from \mathbb{N} to \mathbb{C} for which there exist $a, \chi \in \mathbb{R}$ such that $0 < \chi < 1$ and $|\phi(n)| \leq a\chi^n$. There exists an effective bound m such that if*

$$\lambda^n = \phi(n) + b \tag{12}$$

then $n < m$. Moreover, if $\lambda, b \in \mathbb{A}$ and $a, \chi \in \mathbb{Q}$ are given as input, then m is at most exponential in the length of the input $L = \|\lambda\| + \|b\| + \|a\| + \|\chi\|$.

Proof. The left-hand side of (12) describes points on the unit circle, whereas the right-hand side tends to b . If $|b| \neq 1$, then for n large enough, the right-hand side of (12) will always be off the unit circle. This happens when

$$n > \frac{\ln(|b| - 1/a)}{\ln(\chi)}$$

The difficult case is when b is on the unit circle. We will use Baker's theorem to derive a bound on n . Consider the angle Λ between λ^n and b . This angle can

be zero for at most one value of n , because λ is not a root of unity. Otherwise, we have

$$A = \ln \frac{\lambda^n}{b} = n \ln(\lambda) - \ln b + 2k_n \ln(-1) \neq 0$$

where k_n is an integer chosen so that $A = i\tau$ for some $\tau \in [0, 2\pi)$. Then $2n$ is a height bound for the coefficients in front of the logarithms (because $k_n \leq n$), $H = \max\{H_\lambda, H_b, 4\}$ is a height bound for the arguments to the logarithms and $d = \max\{n_\lambda, n_b\}$ is a bound on the degrees. Then by Baker's theorem, we have

$$\ln |A| > -(48d)^{600} \ln^2 H \ln(2n \ln^2 H)$$

which is equivalent to

$$|A| > (2n \ln^2 H)^{-(48d)^{600} \ln^2 H}$$

This is a lower bound on the length of the arc between λ^n and b . The length of the chord is at least half of the bound: $|\lambda^n - b| \geq |A|/2$. So in the equation $\lambda^n - b = \phi(n)$, the left-hand side is bounded below by an inverse polynomial in n . However, the right-hand side shrinks exponentially quickly. For n large enough, the right-hand side will forever be smaller in magnitude than the left-hand side.

We will now quantify the bound on n . Let $p_1 = (48d)^{600} \ln^2 H$ and $p_2 = 2 \ln^2 H$. Observe that if λ and b are canonically represented algebraic numbers, then p_1, p_2 are polynomials in the size of the input. Then (12) cannot hold if

$$\frac{1}{2} (p_2 n)^{-p_1} \geq a \chi^n$$

which is equivalent to

$$-\ln(2) - \ln(a) - p_1 \ln(p_2) - p_1 \ln(n) \geq n \ln(\chi)$$

Define $p_3 = \ln(2) + \ln(a) + p_1 \ln(p_2)$ and $p_4 = \max\{p_3, p_1\}$ (also polynomials in the size of the input). Then it suffices to have

$$\frac{p_4}{-\ln(\chi)} \leq \frac{n}{1 + \ln(n)}$$

which is guaranteed by

$$\sqrt{n} \geq \frac{p_4}{-\ln(\chi)}$$

Observe that $-1/\ln(\chi)$ is at most exponentially large in $\|\chi\|$. Therefore, the bound on n is exponential in the size of the input. \square

Lemma 4. Suppose $\lambda_1, \lambda_2, a, b, c \in \mathbb{C}$ are non-zero, where $|\lambda_1| = |\lambda_2| = 1$ and λ_1, λ_2 are not roots of unity. Let $\phi(n)$ be a function from \mathbb{N} to \mathbb{C} such that $0 < |\phi(n)| \leq w\chi^n$ for some $w, \chi \in \mathbb{R}$, $\chi \in (0, 1)$. Then there exists a computable bound m such that if

$$a\lambda_1^n = b\lambda_2^n + c + \phi(n) \tag{13}$$

then $n < m$. Moreover, if $\lambda_1, \lambda_2, a, b, c \in \mathbb{A}$ and $w, \chi \in \mathbb{Q}$ are given, then m is at most exponentially large in the length of the input $\|\lambda_1\| + \|\lambda_2\| + \|a\| + \|b\| + \|c\| + \|w\| + \|\chi\|$.

Proof. Multiplying the equation by $\bar{c}/|c||a|$ allows us to assume that $|a| = 1$ and $c \in \mathbb{R}^+$.


Let $f(n) = a\lambda_1^n$, $g(n) = b\lambda_2^n + c$. It is clear that $f(n)$ describes points on the unit circle \mathcal{O}_1 , whilst $g(n)$ describes points on the circle \mathcal{O}_2 with centre c on the real line and radius $|b|$.

If these circles do not intersect, then for n large enough, $|\phi(n)|$ will be forever smaller than the smallest distance between the circles. This happens when

$$n > \frac{\ln(c - |b| - 1) - \ln(w)}{\ln(\chi)}$$

which is an exponential bound in the size of the input.

Suppose now the circles intersect in two points, z_1 and z_2 . Let L_1 be the horizontal line through z_1 and L_2 the horizontal line through z_2 . Let $L_1 \cap \mathcal{O}_1 = \{x_1, z_1\}$, $L_1 \cap \mathcal{O}_2 = \{y_1, z_1\}$, $L_2 \cap \mathcal{O}_1 = \{x_2, z_2\}$ and $L_2 \cap \mathcal{O}_2 = \{y_2, z_2\}$. It is trivial that $z_2 = \bar{z}_1$, $x_2 = \bar{x}_1$, $y_2 = \bar{y}_1$.



We first argue that for n large enough, (13) can hold only if for some intersection point z_i , $\operatorname{Re}(z_i)$ lies between $\operatorname{Re}(f(n))$ and $\operatorname{Re}(g(n))$, or $\operatorname{Im}(z_i)$ lies between $\operatorname{Im}(f(n))$ and $\operatorname{Im}(g(n))$. This can only be violated in two symmetric situations: either $f(n)$ is on the arc z_1z_2 of \mathcal{O}_1 which lies inside \mathcal{O}_2 and $g(n)$ is on the arc y_1y_2 of \mathcal{O}_2 which lies outside \mathcal{O}_1 , or $f(n)$ is on the arc x_1x_2 of \mathcal{O}_1 which lies outside \mathcal{O}_2 and $g(n)$ is on the arc z_1z_2 of \mathcal{O}_2 which lies inside \mathcal{O}_1 . In the first situation, when $g(n)$ is on the arc y_1y_2 of \mathcal{O}_2 outside \mathcal{O}_1 , we have

$$|f(n) - g(n)| \geq |g(n)| - 1 \geq |y_1| - 1$$

This lower bound is positive and independent of n , so equality cannot hold for n large enough because $\phi(n)$ tends to zero exponentially quickly. This is the case when

$$n > \frac{\ln(|y_1| - 1) - \ln(w)}{\ln(\chi)}$$

which is exponentially large in the size of the input. The second situation is analogous.

Therefore, we can assume that one of the intersection points z_i separates $f(n)$ and $g(n)$ horizontally or vertically. We will show a lower bound on $|f(n) - g(n)|$ which shrinks slower than exponentially. The horizontal and vertical cases are completely analogous. We show the working for the horizontal case. Assume that $Re(z_i)$ lies between $Re(f(n))$ and $Re(g(n))$. Clearly,

$$|f(n) - g(n)| \geq |Re(g(n) - f(n))| = |Re(z_i - f(n))| + |Re(g(n) - z_i)|$$

Let $\alpha = \arg(\lambda_1)$, $\gamma = \arg(a)$ and $\beta = \arg(z_i)$. Then

$$|Re(z_i - f(n))| = |\cos(n\alpha + \gamma) - \cos(\beta)| = 2 \left| \sin \frac{\beta - n\alpha - \gamma}{2} \sin \frac{\beta + n\alpha + \gamma}{2} \right|$$

Let u_n, v_n be appropriately chosen integers so that

$$\begin{aligned} \frac{\beta - n\alpha - \gamma}{2} + u_n\pi &\in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right] \\ \frac{\beta + n\alpha + \gamma}{2} + v_n\pi &\in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right] \end{aligned}$$

Then using the inequality

$$|\sin(x)| \geq \frac{|x|}{\pi} \text{ for } x \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$$

we have

$$\begin{aligned} \left| \sin \frac{\beta - n\alpha - \gamma}{2} \right| &\geq \frac{1}{\pi} \left| \frac{\beta - n\alpha - \gamma}{2} + \pi u_n \right| \\ \left| \sin \frac{\beta + n\alpha + \gamma}{2} \right| &\geq \frac{1}{\pi} \left| \frac{\beta + n\alpha + \gamma}{2} + \pi v_n \right| \end{aligned}$$

Both of these expressions are sums of logarithms of algebraic numbers, so we can give lower bounds for them using Baker's theorem as in Lemma 3:

$$|Re(z_i - f(n))| \geq (p_1 n)^{-p_2}$$

for some constants $p_1, p_2 > 0$ which are polynomially large in the input. A similar lower bound holds for $|Re(g(n) - z_i)|$. If $\delta = \arg(\lambda_2)$, $\eta = \arg(b)$ and $\theta = \arg(z_i - c)$, we have

$$|Re(g(n) - z_i)| = |b| (\cos(n\delta + \eta) - \cos(\theta)) \geq (p_3 n)^{-p_4}$$

where p_3, p_4 are positive constants of polynomial size in the input. Hence we have

$$|f(n) - g(n)| \geq 2(p_5 n)^{-p_6}$$

where $p_5 = \max\{p_1, p_3\}$ and $p_6 = \max\{p_2, p_4\}$. Since $\phi(n)$ shrinks exponentially quickly, a bound on n follows past which (13) cannot hold. In the manner of

Lemma 3, we can show that this bound is exponentially large in the size of the input. The vertical case is analogous, except that considering imaginary parts gives sines instead of cosines, so we shift all angles by $\pi/2$ and proceed as above. If the circles are tangent and neither lies inside the other, then the intersection point separates $f(n)$ and $g(n)$ horizontally, so we are done by the above analysis.

Finally, suppose that the circles are tangent and one lies inside the other: $|b| + c = 1$. The argument of $f(n)$ is $\gamma + n\alpha$. By the cosine theorem applied to the triangle with vertices $f(n)$ and the centres of the circles, we have

$$|f(n) - c|^2 = c^2 + 1 - 2c \cos(\gamma + n\alpha)$$

Therefore, the shortest distance from $f(n)$ to a point on \mathcal{O}_2 is

$$h(n) = \sqrt{c^2 + 1 - 2c \cos(\gamma + n\alpha)} - (1 - c)$$

Let $A(n) = \sqrt{c^2 + 1 - 2c \cos(\gamma + n\alpha)}$ and $B = 1 - c$. Since $A \leq 1 + c$, we have $A + B \leq 2$, so

$$h(n) = A - B = \frac{A^2 - B^2}{A + B} \geq c(1 - \cos(\gamma + n\alpha))$$

Let k_n be an integer, so that

$$\gamma + n\alpha + k_n 2\pi \in [-\pi, \pi)$$

A lower bound on this angle follows from Baker's theorem:

$$|\gamma + n\alpha + k_n 2\pi| \geq (p_7 n)^{-p_8}$$

for some constants $p_7, p_8 > 0$ which are polynomially large in the input. Then

$$\cos(\gamma + n\alpha) \leq \cos\left((p_7 n)^{-p_8}\right)$$

so

$$h(n) \geq c \left(1 - \cos\left((p_7 n)^{-p_8}\right)\right)$$

From the Taylor expansion of $\cos(x)$, it follows easily that

$$1 - \cos(x) \geq \frac{11}{24} x^2 \text{ for } x \leq 1$$

Since $p_7, p_8 \geq 1$, we have $(p_7 n)^{-p_8} \leq 1$. Therefore,

$$h(n) \geq c \frac{11}{24} (p_7 n)^{-2p_8}$$

This lower bound on $h(n)$ shrinks inverse-polynomially as n grows. Recall that $h(n)$ is the smallest distance from $f(n)$ to \mathcal{O}_2 . It follows that for n large enough, $|\phi(n)| < h(n)$ forever, so $f(n) = g(n) + \phi(n)$ cannot hold. In the manner of Lemma 3, we can show that the bound is exponentially large in the input. \square

Theorem 2. (van der Poorten [21]) Let $\alpha_1, \dots, \alpha_n$ be algebraic numbers of degree at most d belonging to a number field \mathbb{K} and with heights respectively not exceeding A_1, \dots, A_n . Let P be a prime ideal of \mathbb{K} containing the rational prime p . The following inequalities have no solutions in rational integers b_1, \dots, b_n with absolute values at most $B \geq e^2$:

$$\infty > v_P \left(\alpha_1^{b_1} \dots \alpha_n^{b_n} - 1 \right) > (16(n+1)d)^{12(n+1)} (p^d / \ln(p)) \Omega(\ln(B))^2$$

where $\Omega = \ln(A_1) \dots \ln(A_n)$.

D Skolem's problem, depth 2

In this section, we consider the problem of finding whether a linear recurrent sequence u_n of depth 2 contains 0 as an element. The characteristic equation of the recurrence may have one repeated root $\theta \neq 0$, or two distinct roots $\theta_{1,2}$, giving either

$$u_n = (A + Bn)\theta^n$$

or

$$u_n = A\theta_1^n + B\theta_2^n$$

Solving the problem in the former case is trivial. The latter case is an instance of the *algebraic number power problem*: decide whether there exists $n \in \mathbb{N}$ such that

$$\alpha^n = \beta \tag{14}$$

for given $\alpha, \beta \in \mathbb{A}$. The algebraic number power problem is decidable [12]. Kannan and Lipton [1] proved a polynomial bound on n when β has the form $p(\alpha)$ for a given $p \in \mathbb{Q}[x]$ and α is not a root of unity. We give a brief recapitulation of the decidability proof and extract a polynomial bound on n from it.

Lemma 5. Suppose $\alpha, \beta \in \mathbb{A}$. If α is not a root of unity, then there exists a computable bound m such that if (14) holds, then $n < m$. Moreover, m is polynomial in the length of the input $\|\alpha\| + \|\beta\|$.

Proof. Let $\mathbb{K} = \mathbb{Q}(\alpha, \beta)$. If α is not an algebraic integer, then by Lemma 2 there exists a prime ideal P in the ring $\mathbb{K}_{\mathbb{Z}}$ such that $v_P(\alpha) \neq 0$. Then if $\alpha^n = \beta$, we have

$$v_P(\alpha^n) = nv_P(\alpha) = v_P(\beta)$$

If $v_P(\alpha)$ and $v_P(\beta)$ have different signs, then we are done. Otherwise,

$$n = \frac{v_P(\beta)}{v_P(\alpha)} \leq |v_P(\beta)| \leq \log_2 |\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\beta)| \leq \log_2 |\mathcal{N}_{abs}(\beta)|^d$$

where $d = [\mathbb{Q}(\alpha, \beta) : \mathbb{Q}]$ is at most polynomially large in $\|\alpha\| + \|\beta\|$. It follows that the bound on n is polynomially large in the length of the input.

Suppose α is an algebraic integer. It is not a root of unity, so by Kronecker's theorem [22], α has a Galois conjugate $\sigma(\alpha)$ with magnitude strictly greater than 1. In fact, a significant strengthening of Kronecker's theorem, due to Blanksby and Montgomery [23], guarantees the existence of a conjugate $\sigma(\alpha)$ such that

$$|\sigma(\alpha)| > 1 + \frac{1}{30n_\alpha^2 \ln(6n_\alpha)}$$

which implies

$$\frac{1}{\ln |\sigma(\alpha)|} < 60n_\alpha^2 \ln(6n_\alpha)$$

Then if $\alpha^n = \beta$, we have

$$n = \frac{\ln |\sigma(\beta)|}{\ln |\sigma(\alpha)|} < \ln |\sigma(\beta)| 60n_\alpha^2 \ln(6n_\alpha)$$

Observe that if we are given a canonical description of β , both $\ln |\sigma(\beta)|$ is at most polynomially large in $\|\beta\|$, whereas $60n_\alpha^2 \ln(6n_\alpha)$ is at most polynomially large in $\|\alpha\|$. It follows that the bound on n is polynomial in the length of the input. \square

The condition that α is not a root of unity is obviously necessary in Lemma 5, because if β is also a root of unity, $\alpha^n = \beta$ could hold for all n in a linear set. It is easy to exhibit linear recurrences of depth 2 with infinitely many zeroes (for example, $u_1 = 0$, $u_2 = 1$, $u_{n+2} = u_n$), but by the Lemma, they all have two distinct roots whose ratio is a root of unity.

E Skolem's problem, depth 3

In this section we will focus on Skolem's problem for linear recurrent sequences of depth 3. The characteristic equation of such a sequence may have either three distinct roots α, β, γ , or one repeated root α and one simple root β . If the three roots are distinct, we are concerned with solving for $n \in \mathbb{N}$ equations of the form

$$A\alpha^n + B\beta^n + C\gamma^n = 0 \tag{15}$$

where $A, B, C, \alpha, \beta, \gamma \in \mathbb{A}$ are given and non-zero (if any of them is 0, then the sequence satisfies a recurrence relation of smaller depth). Then (15) is equivalent to

$$\left(\frac{\beta}{\alpha}\right)^n = -\frac{C}{B} \left(\frac{\gamma}{\alpha}\right)^n - \frac{A}{B} \tag{16}$$

We will consider only non-singular sequences: the ratios of the roots α, β, γ are not roots of unity. Let also $|\alpha| \geq |\beta| \geq |\gamma|$. In Lemmas 6, 7, 8 below, the length of the input is $\|A\| + \|B\| + \|C\| + \|\alpha\| + \|\beta\| + \|\gamma\|$.

Lemma 6. *If $|\alpha| > |\beta|$, then there exists an effective bound m such that if equation (15) holds, then $n < m$. Moreover, m is at most exponential in the length of input.*

Proof. Trivial, because $A\alpha^n$ is dominant. If

$$n > \max \left\{ \frac{\ln |A/2B|}{\ln |\beta/\alpha|}, \frac{\ln |A/2C|}{\ln |\gamma/\alpha|} \right\}$$

then

$$\left| -\frac{B}{A} \left(\frac{\beta}{\alpha} \right)^n - \frac{C}{A} \left(\frac{\gamma}{\alpha} \right)^n \right| \leq \left| \frac{B}{A} \left(\frac{\beta}{\alpha} \right)^n \right| + \left| \frac{C}{A} \left(\frac{\gamma}{\alpha} \right)^n \right| < \frac{1}{2} + \frac{1}{2} = 1$$

□

Lemma 7. *If $|\alpha| = |\beta| > |\gamma|$, then there exists an effective bound m such that if equation (15) holds, then $n < m$. Moreover, m is at most exponential in the length of the input.*

Proof. This is a direct application of Lemma 3 to equation (16). □

Lemma 8. *If $|\alpha| = |\beta| = |\gamma|$, then there exist at most two values of n such that equation (15) holds. Moreover, they are at most exponential in the length of the input and are computable in polynomial time.*

Proof. The left-hand side of (16) as a function of n describes points on the unit circle in the complex plane, whereas the right-hand side describes points on a circle centred at $-A/B$ with radius $|C/B|$. Note these circles do not coincide, because $A \neq 0$. We can obtain their equations and compute their intersection point(s). If they do not intersect, then equation (15) can never hold. Otherwise, the equation can only hold if the two sides are simultaneously equal to the same intersection point. For each of the (at most two) intersection points θ , let

$$S_1 = \left\{ n \mid \left(\frac{\beta}{\alpha} \right)^n = \theta \right\}$$

$$S_2 = \left\{ n \mid -\frac{C}{B} \left(\frac{\gamma}{\alpha} \right)^n - \frac{A}{B} = \theta \right\}$$

Observe that $|S_i| \leq 1$, because β/α and γ/α are not roots of unity. We compute S_1 and S_2 from the bound in Lemma 5 and check whether $S_1 \cap S_2$ is non-empty.

□

Next, we consider recurrent sequences of depth 3 with one repeated and one simple root. We are given non-zero $A, B, C, \alpha, \beta \in \mathbb{A}$, and the length of the input is $\|A\| + \|B\| + \|C\| + \|\alpha\| + \|\beta\|$. We wish to solve for n

$$(A + Bn)\alpha^n + C\beta^n = 0 \tag{17}$$

Lemma 9. *There exists an effective bound m such that if (9) holds, then $n < m$. Moreover, m is at most exponential in the length of the input.*

Proof. If $|\alpha| \geq |\beta|$, then for

$$n > \frac{|A| + |C|}{|B|}$$

we have

$$|C| < |B|n - |A| \leq |A + Bn|$$

so

$$|C\beta^n| < |(A + Bn)\alpha^n|$$

and 17 cannot hold. Now suppose $|\alpha| \geq |\beta|$ and rewrite (17) as

$$\frac{A + Bn}{C} = -\left(\frac{\beta}{\alpha}\right)^n$$

A bound on n follows from

$$\left|\frac{A}{C}\right| + \left|\frac{B}{C}\right|n < \left|\frac{\beta}{\alpha}\right|^n$$

which is implied by

$$d(n+1) < \left|\frac{\beta}{\alpha}\right|^n$$

where $d = \max\{|A/C|, |B/C|\}$. Taking logarithms, we see that it suffices to have

$$\frac{n}{1 + \ln(n+1)} > \frac{f}{\ln(\beta/\alpha)}$$

where $f = \max\{\ln(d), 1\}$. Noting that $1 + \ln(n+1) < 2\sqrt{n}$ for all $n \geq 1$, we see that it suffices to have

$$n > 4f^2 / \ln^2(\beta/\alpha)$$

This is an exponential bound in the length of the input. \square

F Skolem's problem, depth 4

In this section we will give lemmas which form a decidability proof for Skolem's problem for depth 4. Assume algebraic numbers A, B, C, D and $\alpha, \beta, \gamma, \delta$ are given and the input has length $\|A\| + \|B\| + \|C\| + \|D\| + \|\alpha\| + \|\beta\| + \|\gamma\| + \|\delta\|$. We wish to solve for n the following equations:

$$A\alpha^n + B\beta^n + C\gamma^n + D\delta^n = 0 \tag{18}$$

$$(A + Bn)\alpha^n + C\beta^n + D\gamma^n = 0 \tag{19}$$

$$(A + Bn)\alpha^n + (C + Dn)\beta^n = 0 \tag{20}$$

$$(A + Bn + Cn^2)\alpha^n + D\beta^n = 0 \tag{21}$$

$$(A + Bn + Cn^2 + Dn^3)\alpha^n = 0 \tag{22}$$

As before, we assume that the ratios of $\alpha, \beta, \gamma, \delta$ are not roots of unity. Solving (22) is trivial. We can rearrange (21) as

$$(A + Bn + Cn^2) \left(\frac{\alpha}{\beta} \right)^n = -D \quad (23)$$

The left-hand side tends to 0 or ∞ in magnitude, depending on whether $|\alpha| < |\beta|$ or not. In both cases, a computable bound on n follows. The remaining equations (18)(19)(20) are more involved.

Lemma 10. *There exists an effective bound m such that if equation (20) holds, then $n < m$. Moreover, this bound is at most exponential in the size of the input.*

Proof. Rearrange (20) as

$$\lambda^n = -\frac{(C + Dn)}{(A + Bn)} \quad (24)$$

where $\lambda = \alpha/\beta$. The right-hand side of (24) tends to $-D/B$ as n tends to infinity.

First, if λ is an algebraic integer, then by Blanksby and Montgomery's theorem [23], it has a Galois conjugate $\sigma(\lambda)$ such that

$$|\sigma(\lambda)| > 1 + \frac{1}{30n_\lambda^2 \ln(6n_\lambda)}$$

We apply the monomorphism σ to both sides of (24), so that the right-hand side tends to $\sigma(D/B)$ whereas the left increases in magnitude exponentially quickly, giving a bound on n .

Second, suppose λ is not an algebraic integer. Then by Lemma 2 there exists a prime ideal P in the ring of integers of $\mathbb{K} = \mathbb{Q}(\alpha, \beta, A, B, C, D)$ such that $v_P(\lambda) \neq 0$. Without loss of generality, we can assume $v_P(\lambda) > 0$ (if $v_P(\lambda) < 0$, swap α with β , A with C , and B with D). Applying v_P to (24) gives

$$v_P(\lambda^n) = nv_P(\lambda) = v_P\left(-\frac{C + Dn}{A + Bn}\right) \leq \ln \left| \mathcal{N}_{\mathbb{K}/\mathbb{Q}}\left(-\frac{C + Dn}{A + Bn}\right) \right|$$

We will now show that for $n > e_1$ we have $\ln \left| \mathcal{N}_{\mathbb{K}/\mathbb{Q}}\left(-\frac{C + Dn}{A + Bn}\right) \right| < p_1$, where e_1 and p_1 are an exponential and a polynomial constant in the input, respectively. This suffices because $nv_P(\lambda) \geq n$ grows with n .

Consider

$$f(n) = \left| \frac{C + Dn}{A + Bn} \right|$$

By the triangle inequality, we have

$$h(n) = \frac{|D|n - |C|}{|B|n + |A|} \leq f(n) \leq \frac{|D|n + |C|}{|B|n - |A|} = g(n)$$

Assume that $n > |A/B|$. It is easy to see that h and g tend to $|D/B|$, h from below and g from above. For

$$n \geq \frac{2|BC| + 2|AD| + |AB|}{|B|^2}$$

we have

$$g(n) \leq \left\lfloor \frac{D}{B} \right\rfloor + \frac{1}{2}$$

Similarly, for

$$n \geq \frac{2|BC| + 2|AD| - |AB|}{|B|^2}$$

we have

$$h(n) \geq \left\lfloor \frac{D}{B} \right\rfloor + \frac{1}{2}$$

Thus, for n at most exponentially large, we have $f(n) \leq |D/B| + 1$. This bound is at most exponential in the size of the input.

Applying the monomorphisms σ from \mathbb{K} into \mathbb{C} to

$$-\frac{C + Dn}{A + Bn}$$

gives terms of the form

$$-\frac{\sigma(C) + \sigma(D)n}{\sigma(A) + \sigma(B)n}$$

We can obtain an upper bound, exponential in the input and independent of n , on the magnitude of each such term. There are exactly $[\mathbb{K} : \mathbb{Q}]$ such terms, one for each monomorphism σ of \mathbb{K} into \mathbb{C} . This gives a polynomially large upper bound on $\ln \left| \mathcal{N}_{\mathbb{K}/\mathbb{Q}} \left(-\frac{C+Dn}{A+Bn} \right) \right|$ which holds for exponentially large n . \square

Lemma 11. *There exists an effective bound m such that if equation (19) holds, then $n < m$. This bound is at most exponential in the size of the input.*

Proof. First suppose $|\alpha| \geq |\beta|, |\gamma|$. Then the term $(A + Bn)\alpha^n$ is dominant. More precisely, rewrite (19) as

$$A + Bn = -C \left(\frac{\beta}{\alpha} \right)^n - D \left(\frac{\gamma}{\alpha} \right)^n$$

and observe that if

$$n > \frac{|A| + |C| + |D|}{|B|}$$

then

$$|A + Bn| \geq |B|n - |A| > |C| + |D| \geq \left| -C \left(\frac{\beta}{\alpha} \right)^n - D \left(\frac{\gamma}{\alpha} \right)^n \right|$$

so (19) cannot hold.

Second, suppose that $|\beta| > |\alpha|, |\gamma|$. Then the term $C\beta^n$ is dominant. More precisely, rewrite (19) as

$$(A + Bn) \left(\frac{\alpha}{\beta} \right)^n + D \left(\frac{\gamma}{\beta} \right)^n = -C$$

We will require that

$$\left| D \left(\frac{\gamma}{\beta} \right)^n \right| < \frac{|C|}{2}$$

and

$$\left| (A + Bn) \left(\frac{\alpha}{\beta} \right)^n \right| < \frac{|C|}{2}$$

The former inequality holds for $n > \ln |C/2D| / \ln |\gamma/\beta|$, which is at most exponentially large in the input. The latter inequality is implied by

$$\left| (n+1) \left(\frac{\alpha}{\beta} \right)^n \right| < \frac{|C|}{2M}$$

where $M = \max \{|A|, |B|\}$. Now let $r = \lceil -\ln(2) / \ln(\alpha/\beta) \rceil$, so that

$$\left(\frac{\alpha}{\beta} \right)^r < \frac{1}{2}$$

and consider only n of the form $n = kr$ for $k \in \mathbb{Z}^+$. If

$$k > \frac{\ln |C/4Mr|}{\ln(7/8)}$$

and $k \geq 5$, we have

$$\left(\frac{\alpha}{\beta} \right)^{kr} k < \left(\frac{1}{2} \right)^k (k+1) < \left(\frac{7}{8} \right)^k < \frac{|C|}{4Mr}$$

so

$$\left(\frac{\alpha}{\beta} \right)^n (n+1) \leq \left(\frac{\alpha}{\beta} \right)^n 2n < \frac{|C|}{2M}$$

It is clear that r is at most exponentially large in the size of the input, whereas the bound on k is polynomial. Therefore, the bound on n is exponential.

Finally, suppose $|\beta| = |\gamma| > |\alpha|$. Rewrite (19) as

$$\left(\frac{\beta}{\gamma} \right)^n = -\frac{D}{C} - \frac{A + Bn}{C} \left(\frac{\alpha}{\gamma} \right)^n$$

Then an exponential bound on n follows from Lemma 3, because the right-hand side is a constant plus an exponentially decaying term, whereas the left-hand side is on unit circle. \square

Lemma 12. *If $\alpha, \beta, \gamma, \delta$ do not all have the same magnitude, then there exists an effective bound m such that if equation (18) holds, then $n < m$. Moreover, m is at most exponentially large in the input size.*

Proof. Let $|\alpha| \geq |\beta| \geq |\gamma| \geq |\delta|$. First, if $|\alpha| > |\beta|$, then $A\alpha^n$ is the dominant term in (18). Rewrite the equation as

$$\frac{B}{A} \left(\frac{\beta}{\alpha} \right)^n + \frac{C}{A} \left(\frac{\gamma}{\alpha} \right)^n + \frac{D}{A} \left(\frac{\delta}{\alpha} \right)^n = -1$$

and observe that if

$$n > \max \left\{ \frac{\ln |3B/A|}{\ln |\alpha/\beta|}, \frac{\ln |3C/A|}{\ln |\alpha/\gamma|}, \frac{\ln |3D/A|}{\ln |\alpha/\delta|} \right\}$$

then

$$\left| \frac{B}{A} \left(\frac{\beta}{\alpha} \right)^n + \frac{C}{A} \left(\frac{\gamma}{\alpha} \right)^n + \frac{D}{A} \left(\frac{\delta}{\alpha} \right)^n \right| < \frac{1}{3} + \frac{1}{3} + \frac{1}{3} = 1$$

Second, if $|\alpha| = |\beta| > |\gamma|$, then rewrite (18) as

$$\left(\frac{\beta}{\alpha} \right)^n = -\frac{A}{B} - \frac{C}{B} \left(\frac{\gamma}{\alpha} \right)^n - \frac{D}{B} \left(\frac{\delta}{\alpha} \right)^n \quad (25)$$

The left-hand side of (25) is on the unit circle, whereas the right is a constant plus exponentially decaying terms. An exponential bound on n follows from Lemma 3.

Finally, if $|\alpha| = |\beta| = |\gamma| > |\delta|$, then an exponential bound on n follows from Lemma 4 applied to equation (25). \square

Thus, the only outstanding problem is to solve (18) when $|\alpha| = |\beta| = |\gamma| = |\delta|$. This is difficult for general $\alpha, \beta, \gamma, \delta$, so we will restrict ourselves to two sufficient special cases: when at least two of them are real, and when they are two pairs of complex conjugates. We will also assume that they are *algebraic integers*. This is sufficient for our application to the orbit problem, because any orbit instance $\exists n. A^n x \in V$ may be reduced in polynomial time to an instance where A is an integer matrix, so that its eigenvalues are algebraic integers. This trick is not readily available when solving a Skolem instance.

Lemma 13. *Let $\alpha, \beta, \gamma, \delta$ be algebraic integers with $|\alpha| = |\beta| = |\gamma| = |\delta|$. If $\alpha, \beta \in \mathbb{R}$ or (α, β) and (γ, δ) are pairwise complex conjugates, there exists an effective bound m such that if equation (18) holds, then $n < m$. Moreover, m is exponential in the length of the input.*

Proof. Let $\mathbb{K} = \mathbb{Q}(\alpha, \beta, \gamma, \delta, A, B, C, D)$. First suppose that $\alpha, \beta \in \mathbb{R}$. If $\alpha = \beta$, then we have a Skolem instance of depth 3

$$\exists n. (A + B)\alpha^n + C\gamma^n + D\gamma^n = 0$$

An exponential bound on n follows from Lemma 8. If $\alpha = -\beta$, then we consider even n and odd n separately, obtaining two Skolem instances of depth 3 of the same type.

Now suppose $\beta = \bar{\alpha}$ and $\gamma = \bar{\delta}$. If α/β is an algebraic integer, then since it is not a root of unity, there exists a monomorphism σ from \mathbb{K} to \mathbb{C} such that

$|\sigma(\alpha)| \neq |\sigma(\beta)|$. Applying σ to (18) leads to a Skolem instance of depth 4 with roots whose magnitudes are not all the same. A bound on n follows from Lemma 12.

Suppose then that α/β is not an algebraic integer. By the reasoning of Lemma 2, there exists a prime ideal P in $\mathbb{K}_{\mathbb{Z}}$ such that $v_P(\alpha) \neq v_P(\beta)$ and at least one of $v_P(\alpha)$ and $v_P(\beta)$ is strictly positive. Assume without loss of generality that

$$v_P(\alpha) > v_P(\beta) \geq 0$$

Since $\alpha\beta = \gamma\delta = |\alpha|^2$, we have

$$v_P(\alpha) + v_P(\beta) = v_P(\gamma) + v_P(\delta)$$

Therefore, at most two of the roots are smallest under the valuation v_P .

If one root, say β , is strictly smaller under v_P than the rest, then rewrite (18) as

$$A\alpha^n + B\beta^n = -C\gamma^n - D\delta^n \quad (26)$$

Since $v_P(\beta) < v_P(\alpha)$, for $n > v_P(A/B)/v_P(\beta/\alpha)$ we have

$$v_P(A\alpha^n + B\beta^n) = v_P(B) + nv_P(\beta)$$

whereas

$$v_P(-C\gamma^n - D\delta^n) \geq v_P(C) + nv_P(\gamma)$$

Therefore, for $n > v_P(B/C)/v_P(\gamma/\beta)$, we have that the left-hand side of (26) is strictly smaller under v_P than the right-hand side, so (18) cannot hold. This bound on n is polynomial in the input size.

Now suppose that there are two roots with strictly smallest valuation:

$$0 \leq v_P(\beta) = v_P(\gamma) < v_P(\alpha) = v_P(\delta)$$

Then rewrite (18) as

$$B\beta^n \left(\left(-\frac{C}{B} \right) \left(\frac{\gamma}{\beta} \right)^n - 1 \right) = A\alpha^n + D\delta^n \quad (27)$$

Since γ/β is not a root of unity, the term $(-C/B)(\gamma/\beta)^n - 1$ can be zero for at most one value of n . This value is at most polynomially large in the input size (by Lemma 5). For all other n , we may apply van der Poorten's theorem to this term. Let p be the unique prime rational integer in the ideal P , and let $d = [\mathbb{K} : \mathbb{Q}]$. Let H be an upper bound for the heights of $-C/B$ and γ/β . Then by Theorem 2, we have

$$v_P \left(\left(-\frac{C}{B} \right) \left(\frac{\gamma}{\beta} \right)^n - 1 \right) \leq (48d)^{36} \frac{p^d}{\ln(p)} (\ln H)^2 (\ln n)^2$$

It is classical $\mathcal{N}(P) = p^f$ for some positive integer f , so $\mathcal{N}(P) \geq p$. Moreover, since α is an algebraic integer, all prime ideals P_1, \dots, P_s in the factorisation of $[\alpha]$ appear with positive exponents k_1, \dots, k_s :

$$[\alpha] = P_1^{k_1} \dots P_s^{k_s}$$

Since $\mathcal{N}(P_i) \geq 2$ for all P_i , we have

$$|\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha)| = \mathcal{N}([\alpha]) \geq \mathcal{N}(P) \geq p$$

Therefore, p is at most exponentially large in the input size. Then we can write the inequality from van der Poorten's theorem as

$$v_P \left(\left(-\frac{C}{B} \right) \left(\frac{\gamma}{\beta} \right)^n - 1 \right) \leq E_1 (\ln n)^2$$

where E_1 is exponentially large in the input size and independent of n . Now we apply v_P to both sides of equation (27):

$$v_P(LHS) \leq v_P(B) + nv_P(\beta) + E_1 (\ln n)^2$$

and

$$v_P(RHS) \geq v_P(A) + nv_P(\alpha)$$

Equation (18) cannot hold if

$$v_P(B) + nv_P(\beta) + E_1 (\ln n)^2 < v_P(A) + nv_P(\alpha)$$

which is implied by

$$v_P(B/A) + E_1 (\ln n)^2 < n$$

since $v_P(\alpha) > v_P(\beta)$. Let $E_2 = \max\{v_P(B/A), E_1\}$, then this is implied by

$$E_2 ((\ln n)^2 + 1) < n$$

Since

$$(\ln n)^2 + 1 < \frac{5\sqrt{n}}{2}$$

for all $n \geq 1$, it suffices to have

$$n > \left(\frac{5}{2} E_2 \right)^2$$

This bound on n is exponential in the size of the input. \square

References

1. Kannan, R., Lipton, R.: Polynomial-time algorithm for the orbit problem. *Journal of the ACM* **33**(4) (1986) 808–821
2. Arvind, V., Vijayaraghavan, T.: The orbit problem is in the GapL hierarchy. *J. Comb. Optim.* **21**(1) (2011) 124–137
3. Tarasov, S.P., Vyalyi, M.N.: Orbits of linear maps and regular languages. *CoRR* **abs/1011.1842** (2010)

4. Skolem, T.: Ein Verfahren zur Behandlung gewisser exponentialer Gleichungen und diophantischer Gleichungen. *Skand. Mat. Kongr.* **8** (1934) 163–188
5. Mahler, K.: Eine arithmetische Eigenschaft der Taylor-Koeffizienten rationaler Funktionen. *Proc. Akad. Wet. Amsterdam* **38** (1935) 51–60
6. Lech, C.: A note on recurring series. *Arkiv för Matematik* **2** (1953) 417–421
7. Hansel, G.: Une démonstration simple du théorème de Skolem-Mahler-Lech. *Theoretical Computer Science* **43** (1986) 91 – 98
8. Halava, V.: Decidable and undecidable problems in matrix theory. TUCS Technical Report 127 (1997)
9. Vereshchagin, N.: Occurrence of zero in a linear recursive sequence. *Mathematical Notes* **38** (1985) 609–615
10. Mignotte, M., Shorey, T., Tijdeman, R.: The distance between terms of an algebraic recurrence sequence. *Jour. Reine Angew. Math.* **349** (1984) 63 – 76
11. Baker, A.: *Transcendental number theory*. Cambridge University Press (1975)
12. Halava, V., Harju, T., Hirvensalo, M., Karhumäki, J.: Skolem’s problem – on the border between decidability and undecidability. TUCS Technical Report (683) (2005)
13. Blondel, V., Portier, N.: The presence of a zero in an integer linear recurrent sequence is NP-hard to decide (2002)
14. Schönhage, A.: On the power of random access machines. In Maurer, H., ed.: *Automata, Languages and Programming*. Volume 71 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg (1979) 520–529
15. Mignotte, M.: Some useful bounds. *Computer Algebra* (1982) 259–263
16. Pan, V.: Optimal and nearly optimal algorithms for approximating polynomial zeros. *Computers & Mathematics with Applications* **31**(12) (1996) 97 – 138
17. Cohen, H.: *A Course in Computational Algebraic Number Theory*. Springer (1993)
18. Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. *Mathematische Annalen* **261** (1982) 515–534
19. Zippel, R.: Zero testing of algebraic functions. *Information Processing Letters* **61**(2) (1997) 63 – 67
20. Stewart, I., Tall, D.: *Algebraic Number Theory and Fermat’s Last Theorem*. 3rd edn. A. K. Peters (2002)
21. van der Poorten, A.J.: Linear forms in logarithms in the p-adic case. *Transcendence Theory: Advances and Applications* (1977) 29–57
22. Kronecker, L.: Zwei Sätze über Gleichungen mit ganzzahligen Koeffizienten. *J. Reine Angew. Math.* **53** (1875) 173–175
23. Blanksby, P., Montgomery, H.: Algebraic integers near the unit circle. *Acta Arith.* (1971) 355–369
24. Shidlovskii, A.: *Transcendental Numbers*. New York: de Gruyter studies in mathematics (1989)
25. Kozen, D.: *The Design and Analysis of Algorithms*. Springer-Verlag, New York (1991)
26. Chonev, V.: The orbit problem in zero and one dimensions. Master’s thesis, Oxford University Department of Computer Science (2011)